Charles A. Windom
Hemant Gaidhani

Virtualizing Microsoft Tier 1 Applications with VMware vSphere 4

SYBEX SERIOUS SKILLS.

Virtualizing

Microsoft® Tier 1 Applications with VMware vSphere™ 4

Virtualizing
Microsoft® Tier 1 Applications with VMware vSphere™ 4

Charles A. Windom Sr.

Hemant Gaidhani

Alex Fontana

Wiley Publishing, Inc.

Acquisitions Editor: Agatha Kim
Development Editor: Tom Cirtin
Technical Editor: Shandar Varadarajan
Production Editor: Rachel McConlogue
Copy Editor: Kim Wimpsett
Editorial Manager: Pete Gaughan
Production Manager: Tim Tate
Vice President and Executive Group Publisher: Richard Swadley
Vice President and Publisher: Neil Edde
Book Designers: Maureen Forys and Judy Fung
Compositor: Craig Johnson, Happenstance Type-O-Rama
Proofreader: Publication Services, Inc.
Indexer: Jack Lewis
Project Coordinator, Cover: Lynsey Stanford
Cover Designer: Ryan Sneed
Cover Image: Paul Cooklin / Brand X Pictures / Getty Images

Copyright © 2010 by Wiley Publishing, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN: 978-0-470-56360-1
ISBN: 978-0-470-90814-3 (ebk)
ISBN: 978-0-470-90817-4 (ebk)
ISBN: 978-0-470-90816-7 (ebk)

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at `http://www.wiley.com/go/permissions`.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Web site is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or Web site may provide or recommendations it may make. Further, readers should be aware that Internet Web sites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services or to obtain technical support, please contact our Customer Care Department within the U.S. at (877) 762-2974, outside the U.S. at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Library of Congress Cataloging-in-Publication Data is available from the publisher.

TRADEMARKS: Wiley, the Wiley logo, and the Sybex logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. Microsoft is a registered trademark of Microsoft Corporation. VMware vSphere is a trademark of VMware, Inc. All other trademarks are the property of their respective owners. Wiley Publishing, Inc. is not associated with any product or vendor mentioned in this book.

10 9 8 7 6 5 4 3 2 1

Dear Reader,

Thank you for choosing *Virtualizing Microsoft Tier 1 Applications with VMware vSphere 4*. This book is part of a family of premium-quality Sybex books, all of which are written by outstanding authors who combine practical experience with a gift for teaching.

Sybex was founded in 1976. More than 30 years later, we're still committed to producing consistently exceptional books. With each of our titles, we're working hard to set a new standard for the industry. From the paper we print on to the authors we work with, our goal is to bring you the best books available.

I hope you see all that reflected in these pages. I'd be very interested to hear your comments and get your feedback on how we're doing. Feel free to let me know what you think about this or any other Sybex book by sending me an email at `nedde@wiley.com`. If you think you've found a technical error in this book, please visit `http://sybex.custhelp.com`. Customer feedback is critical to our efforts at Sybex.

Best regards,

Neil Edde
Vice President and Publisher
Sybex, an Imprint of Wiley

To my parents Ossie and Mary Windom; because of your love and guidance, I have been able to realize my greatest dreams. To my beautiful wife Lana, thank you for being the major inspiration in my life and for standing by me when I've had to make decisions that weren't always popular. To my daughter, Shanelle and her newborn son Symari, I am so proud of you. To my boys Charles Jr., Andre and James: you guys are my greatest accomplishments on this earth and I love you all.
—Charles A. Windom Sr.

To my parents, Vasant and Vasudha; whatever I am today is because of your love, encouragement, and upbringing. To my lovely wife, Mrudula, for standing by me during all the ups and downs in our ride together, and to our beautiful daughters, Mitali and Namita, who make each day so joyful and exciting.
—Hemant Gaidhani

To my parents, Alejandro and Yolanda; your dedication and sacrifice has been my inspiration to break free from my comfort zone. To my wife Laura, who has been by my side through this and many adventures, and to Joseph and Sissy—this is ultimately all for you guys.
—Alex Fontana

Acknowledgments

I would like to first thank coauthors Alex and Hemant for their expertise and major contributions to this book. Their constant attention to detail and dedication to this effort made this experience a lot less stressful. Thanks also to Andrea Eubanks and VMware for allowing me to take the virtualization journey. None of this would be possible without them. Thanks also to Jason Miller and Marco Gallo for the excellent support at all of the trade shows. Thanks to the VMware Engineering, Performance, and File System teams for their support. Also thanks to my friend and colleague Shandar Varadarajan for all your valuable feedback and guidance. In addition, thanks to my friends and extended family who supported me through this entire effort. Thanks to Agatha Kim for allowing me to communicate my ideas and thoughts to the world and for providing the much-needed guidance when the waters were murky.

Thanks to Tom Cirtin and the rest of the Sybex team for their excellent suggestions and direction and for riding the team and ensuring things got done. And last but not least, thanks to my mother and father wherever they may be. Rest in peace.

—Charles A. Windom Sr.

I would like to start with special thanks to my coauthors: Charles, who initiated this project and invited me to join in, and Alex, for accepting to join the ride along with us.

I would like to extend a big thank-you to VMware for the opportunity to work on the most exciting virtualization products and be part of this disruptive technology sweeping the IT data centers.

This book would not have been possible without the endless love and support of my family and friends. I would like to thank my parents and parents-in-law for all their love and support. I want to thank my wife, Mrudula, for making sure things were as smooth as usual throughout the writing process and taking care of our daughters, Mitali and Namita.

Thanks to Shandar Varadarajan, our technical editor, for his detailed reviews and for not only making sure we covered the right topics but also ensuring that the content is technically accurate.

Finally, a big thank-you to the entire Sybex team: Agatha Kim, our acquisitions editor, for giving us the opportunity to write this book and accommodating our busy schedules; Tom Cirtin, our developmental editor, for his encouraging words, keeping me on track, and providing excellent feedback at all times; Rachel McConlogue, our production editor; Kim Wimpsett, our copyeditor; and Pete Gaughan and Neil Edde.

—Hemant Gaidhani

Big thanks are in order for Charles and Hemant, my coauthors, for the opportunity to join them in this project. I've enjoyed working with the both of you over the past few months and look forward to future endeavors. Thank you to our technical editors Shandar Varadarajan and Deji Akomolafe; their vast knowledge helped to ensure we delivered the most technically correct content possible. Agatha, Tom, and the entire Sybex team—you guys are awesome. Thank you for keeping us on track and getting this book out the door.

Finally, I would like to give very big thanks for those that have had to put up with the crazy work hours before, during, and beyond this project. Laura, Joe, Sissy, I love you guys, and knowing I get to come home to you all is what gets me through the long hours.

—Alex Fontana

About the Authors

Charles A. Windom, Sr., has more than 30 years of expertise in the computer science field. His experience includes work with designing and deploying Microsoft technologies such as Windows, Active Directory, Exchange Server, SQL Server, and SharePoint for SMB and enterprise corporations. He has worked for Digital Equipment Corporation, Compaq Computer Corporation, and Hitachi Data Systems. He also founded and runs his own company, Intillex Technology Consulting, with Nathan Tran.

Charles is currently a Zimbra solutions architect and former senior Microsoft solutions architect in the VMware Microsoft Solutions Platform Group. He and others on the VMware team are responsible for the testing and qualification of the 64-bit Microsoft Clustering and database availability group support on the VMware Infrastructure v3.5 and vSphere platforms. Charles has presented at several industry events including EMC World, VMworld, Storage Networking World, and Microsoft TechED, and he has authored several solutions documents for VMware about Microsoft applications on the VMware virtualization platform. In his spare time, Charles enjoys spending time with his family, listening to and playing music, and tinkering in his vSphere lab and with the latest technological gadgets.

Hemant Gaidhani has more than 18 years of experience in software development, performance engineering, quality assurance, and technology marketing. Before joining VMware in 2006, Hemant worked at high-technology companies such as EMC, Commerce One, Informix, and Unisys. At VMware, Hemant is a senior technical marketing manager currently focused on VMware's management products. Before that, Hemant specialized in vSphere performance, scalability, and best practices for enterprise applications running in VMware virtualization environments. He has authored numerous publications and has been a regular speaker at industry conferences such as VMworld and EMC World. Hemant holds a bachelor's degree in electrical and electronics engineering from Birla Institute of Technology and Science in Pilani, India, and master's degree in business administration from Santa Clara University.

Hemant lives in Santa Clara, California with his wife, Mrudula, and their daughters, Mitali and Namita. In his spare time, Hemant enjoys listening to music and spending time with his family. He has been teaching software courses at the Santa Clara Adult Education's High Tech Academy.

Alex Fontana is an infrastructure architect at VMware, specializing in Microsoft applications. Alex has worked in the technology industry for more than 10 years and has spent the last 5 years focused on virtualizing Microsoft applications using VMware virtualization products. As a member of the IT Architecture team at VMware, Alex helped design and build the Microsoft Active Directory and Exchange infrastructure still in use today.

Alex lives in the San Francisco Bay Area with his wife, Laura, and two children, Joseph and Sissy. In his spare time Alex enjoys snowboarding, running, biking, and playing golf. Look for Alex presenting at future Microsoft technology conferences and the number-one virtualization event, VMworld.

Contents at a Glance

Introduction . . . . . *xvii*

Chapter 1 • VMware vSphere 4 Overview 1

Chapter 2 • Virtualizing Microsoft Server Applications 41

Chapter 3 • Virtualizing Windows Server 2008 85

Chapter 4 • Virtualizing Active Directory 139

Chapter 5 • Virtualizing Internet Information Services 179

Chapter 6 • Virtualizing Exchange Server 235

Chapter 7 • Virtualizing SQL Server 313

Chapter 8 • Virtualizing Microsoft Office SharePoint Server 2007 371

Chapter 9 • Virtualizing Microsoft Remote Desktop Services. 415

Index *487*

Contents

Introduction . . . . . *xvii*

Chapter 1 • VMware vSphere 4 Overview 1
Infrastructure Services 1
VMware vCompute 2
VMware vStorage 11
VMware vNetwork 15
Application Services 19
Availability 19
Security 24
Scalability 25
Management Services 27
vCenter Server 28
vCenter Orchestrator 33
vCenter Guided Consolidation 33
vCenter Update Manager 34
vSphere Management Assistant 34
vSphere Host Update Utility 35
vSphere Command-Line Interfaces: vCLI and PowerCLI 35
VMware vApp 35
VMware Management Solutions 35
VMware vCenter AppSpeed 36
VMware vCenter CapacityIQ 36
VMware vCenter Chargeback 37
VMware vCenter Lab Manager 37
VMware vCenter Lifecycle Manager 37
VMware vCenter Site Recovery Manager 37
VMware vSphere Editions 38
vSphere for Small Businesses 38
vSphere for Midsize and Enterprise Businesses 39
vSphere Compatibility Guides 39
Summary 40

Chapter 2 • Virtualizing Microsoft Server Applications 41
The New Era of Virtualization 41
Virtualization Technologies 42
Hosted Virtualization 42
Bare-Metal Virtualization 43

Hardware Considerations 44
Planning Virtual Deployment 44
Understanding Operational Requirements 45
Assessing Your IT Infrastructure Using VMware Capacity Planner 46
Licensing Considerations 47
Migrating to a Virtual Environment 48
Initial Creation of Virtual Machines 48
Ongoing Provisioning Strategy 49
Ensuring High Availability 49
Host Availability 50
Application-Aware Availability 50
Host Availability Combined with Application-Aware Availability 50
Devising a Backup Strategy 51
Backing Up Virtual Machines Using VMware Data Recovery 52
Patching Virtual Machines Using VMware Update Manager 65
Deploying VMware Update Manager 65
Using VMware Update Manager 67
Microsoft Server Virtualization Validation Program 78
Deployment Strategies 81
Summary 82

Chapter 3 • Virtualizing Windows Server 2008 85

Understanding Guest OS Compatibility and Requirements 86
Guest OS Compatibility 86
Guest OS Requirements 86
Deploying Windows 2008 Virtual Machines 87
Choosing Virtual Hardware 88
Creating the Virtual Machine 90
Installing Windows 2008 R2 101
Using Templates to Deploy Virtual Machines 108
Converting a Physical Machine to a Virtual Machine 116
Deploying a Windows Failover Cluster in vSphere 118
Clustering Virtual Machines 118
Hardware and Software Requirements for Failover Clusters 119
Building a Microsoft Failover Cluster 120
Protecting Your Data 134
In-Guest Backup Solutions 134
Virtual Machine Backup Solutions 135
Summary 137

Chapter 4 • Virtualizing Active Directory 139

Characterizing Your Active Directory Workload 140
32-bit or 64-bit Active Directory Domain Controllers? 141
64-bit 141
32-bit 141
Network Configuration 142

vSphere Host Placement and Configuration 143
Active Directory Time Synchronization. 144
Kerberos v5 144
Controlling Time Drift. 144
Active Directory Backup and Recovery 147
Securing Your Domain Controller Virtual Machine Files. 149
Monitoring Active Directory Replication. 150
Tools 151
What to Monitor 154
Active Directory Proof of Concept 154
Creating and Configuring the Active Directory Virtual Machine 155
Configure and Synchronize Time 165
Back Up Active Directory Domain Controllers 168
Prevent Active Directory Deletions 171
Monitor Active Directory Replication 173
Summary 176

Chapter 5 • Virtualizing Internet Information Services 179
Deploying Internet Information Services. 180
Choosing a Flavor of Windows Server 2008. 180
Installing IIS 181
IIS Performance 182
Monitoring Performance 184
Performance Tuning. 189
Network Load Balancing. 195
Securing Virtualized Internet Information Services 198
Authentication. 199
Using VMware's vShield Zones Firewall 199
Virtualizing IIS Proof of Concept 201
Installing IIS 202
Configuring ESX for Windows Network Load Balancing 211
Installing VMware AppSpeed 224
Summary 234

Chapter 6 • Virtualizing Exchange Server 235
Exchange Server Virtualization Trends 236
Exchange Server Comparisons 237
Exchange Server 2003. 237
Exchange Server 2007 and 2010 Server Roles 237
Edge Server Role 238
Client Access Server Role 239
Hub Transport Role 241
Mailbox Server Role 242
Unified Messaging Server Role 244

Exchange High Availability 247
Hardware Availability 247
Service Availability. 247
Data Redundancy and Availability 248
Workload Considerations When Virtualizing Exchange Server Roles 250
Edge Transport 250
Client Access Server 255
Hub Transport 258
Client Access and Hub Transport Combined Role 261
Mailbox Server 262
Unified Messaging 270
Backup and Restore 274
VMware Exchange Building Blocks 275
Virtualizing Exchange Server 2010 Proof of Concept 277
Building the Exchange 2010 Server Environment 278
Preparing the Active Directory Environment for Exchange Server 2010 279
Building the Exchange Client Access Server Role Virtual Machine 284
Summary 311

Chapter 7 • Virtualizing SQL Server. 313
SQL Server Virtualization Considerations. 313
Data Integrity in a Virtual Environment 314
SQL Server Consolidation Options 315
SQL Server Performance in a Virtual Environment 316
32-bit vs. 64-bit SQL Server 317
Licensing Considerations 318
Support Considerations 318
Security Considerations. 320
Hardware Considerations 321
Storage Considerations 321
High Availability Considerations 323
Backup Considerations 325
What Not to Virtualize 326
Characterizing SQL Server Deployments for Virtualization 326
Characterizing SQL Server Workloads 327
Capturing Resource Utilization Baseline for Existing Databases. 328
Defining Availability Requirements 331
Defining Backup Strategies and Requirements 331
Defining Patching and Update Strategies. 332
Defining Capacity and System Growth Requirements. 332
Ranking SQL Server Candidates for Virtualization 333
Defining Deployment Strategy 333
Migrate to Virtual Environment. 335

Virtualizing SQL Server Proof of Concept . . . 337
Configuring vSphere . . . 338
Configuring SQL Server Virtual Machines . . . 340
Setting Up SQL Server Failover Clustering . . . 347
Testing Virtual SQL Server Deployment . . . 360
Monitoring Virtual SQL Server Performance . . . 362
Summary . . . 370

Chapter 8 • Virtualizing Microsoft Office SharePoint Server 2007 . . . 371
Why Do You Need SharePoint Governance? . . . 371
What Is SharePoint Governance? . . . 372
SharePoint Server Roles and Responsibilities . . . 374
Considering Server Role Virtualization Workload . . . 375
Web Front-End Server Role (and Query Role) . . . 376
Application Server Role . . . 377
Index Server Role . . . 377
Database Back-End Server Role . . . 378
Sizing the Server Roles . . . 378
Using the Microsoft System Center Capacity Planner 2007 . . . 379
Using the HP Sizing Tool . . . 386
SharePoint Proof of Concept . . . 395
Building and Configuring the SQL Server Database Back-End Virtual Machine . 396
Configuring the SharePoint Central Site Administration Virtual Machine . . . 402
Configuring the SharePoint Index Server Role . . . 404
Configuring the SharePoint Web Front-End Server Role . . . 405
Configuring Windows Network Load Balancing . . . 409
Summary . . . 413

Chapter 9 • Virtualizing Microsoft Remote Desktop Services . . . 415
RDS Virtualization Considerations . . . 415
RD Session Host Role . . . 417
RD Licensing Role . . . 422
RD Connection Broker Role . . . 423
RD Gateway Role . . . 423
RD Web Access Role . . . 424
Characterizing Remote Desktop Services Environment . . . 425
Checking Application Compatibility . . . 427
Defining User Requirements . . . 428
Sizing Your RDS Deployment . . . 430
Deployment Considerations . . . 431
Initial Creation of RDS Virtual Machines . . . 431
Ongoing Provisioning Strategy . . . 432
vSphere DRS and HA Considerations . . . 432

Virtualizing Remote Desktop Services Proof of Concept 433
Configuring the Active Directory Virtual Machine 434
Configuring the RD Licensing Virtual Machine 434
Configuring the RD Session Host Role 439
Configuring the RD Web Access Virtual Machine 450
Configuring the RD Connection Broker Role. 453
Configuring RemoteApp 459
Configuring the RD Gateway Role. 465
Configuring Windows Network Load Balancing 472
Testing Virtual RDS Farm Deployment 480
Monitoring the Virtual RDS Farm 484
Summary 485

Index *487*

Introduction

Over the years as power costs have soared and data center space has become scarce and expensive to acquire, enterprises and small businesses have been looking for ways to decrease their data center footprints as well as reduce the overall costs of operating an IT infrastructure. As the number of different applications being deployed in data centers grows, efficiency in deployment, management, and resource consumption is critical. Imagine having to deploy a new piece of physical hardware for each of these applications—or, even worse, each tier of these applications. Well, that wasn't too long ago. Racking, stacking, and installing the OS took a day to complete. The procurement and approval processes alone took weeks. How did we ever get anything done? Luckily for us, VMware has dramatically changed the landscape of IT with the vSphere suite of virtualization products.

With the introduction of virtualization into the data center, businesses have sought to virtualize the infrastructure of their data centers by deploying hosted and bare-metal virtualization products. It's a slow process for most companies, but, nonetheless, virtualization is extending its reach into data centers everywhere, one virtual machine at a time. Over the past few years, virtualization has gone from being a tool for developers and testers to being the default infrastructure for many organizations' IT departments. As we strive to make our processes more efficient, virtualization just makes sense. Provisioning takes minutes, not days; management can be done through a single pane of glass, and hardware upgrades—rather, virtual-hardware upgrades—happen from home on the weekends.

As virtualization becomes more and more common in today's data centers, businesses are becoming more comfortable deploying virtualization in production. For systems administrators who are new to virtualization, the file, print, DHCP, DNS, and small web servers are the low-hanging fruit. Most systems administrators have been given virtualization initiatives from the head brass, and these "low-risk, low-utilization" applications are the easy wins. As server hardware becomes more powerful and applications become more efficient, though, it makes no sense to deploy a single application on a single piece of physical hardware, even the larger and more demanding applications. These are usually the *tier-1* applications that are considered critical for the business.

When virtualization technologies were first introduced to enterprise and small/medium businesses, many thought this new technology would enable them to deploy all of their infrastructure and business-critical or tier-1 applications on virtualized platforms. When we say *tier-1*, we are talking specifically about infrastructure and applications that are critical to a business's daily operations. Businesses, including enterprises and small/medium ones, would be severely crippled without these applications running optimally. However, early virtualization efforts for these tier-1 applications were not very successful for a variety of reasons. Immaturity of the virtualization platforms as well as improper planning caused these projects to fail and created a negative opinion of virtualization in businesses.

With the introduction of more capable virtualization hardware and new advances in the virtualization platforms, such as the following:

- Advanced/expanded hardware support
- Mature virtual hardware
- Increased performance
- Reduced provisioning times
- Intelligent dynamic load balancing
- Increased uptime for applications
- Expanded physical resources for virtual machines
- Emerging guidance around virtualizing applications

customers are now able to run the most voracious of applications on the vSphere platform. It is these enhancements and advances in virtualization platforms that have customers excited about virtualizing their advanced infrastructures and tier-1, business-critical applications.

It's no secret that Microsoft continues to dominate the operating-system market. That includes the server OS and application realm. According to IDC, Windows-based servers accounted for about 70 percent share of the x86 server revenue as well as server unit shipments in the year 2009.

Look at the headers of any email message in your corporate email inbox, and chances are it will have originated on a Microsoft Exchange server. If you write documentation and use a collaboration or version-control system, it's probably SharePoint. What is at the head and tail ends of that SharePoint environment? Microsoft IIS and Microsoft SQL Server. What do all of these applications ride on top of? We've come full circle and can now understand why the Microsoft operating systems are where they are in the market. It's to your advantage as an IT professional to understand how these Microsoft tier-1 applications fit into the virtual environment and to understand the caveats and the advantages you can gain from deploying on VMware. That is the goal of this book—to help IT professionals gain the confidence and understanding to successfully deploy tier-1 Microsoft applications on vSphere.

Who Should Read This Book

Virtualization administrators responsible for designing and managing the server, storage and network infrastructure deployment for VMware vSphere, will benefit from this book. Understanding the virtualization-specific considerations for tier-1 Microsoft applications will enable them to design optimal architecture that meets necessary service-level agreements in place.

Microsoft Windows and Microsoft Server professionals will relate to the topics covered in this book. For the administrators who manage only a small portion of the Microsoft infrastructure or application stack, this is a good way to learn a little about the virtualized deployment of other tier-1 Microsoft applications. This will undoubtedly add significant value to your career when your colleagues are asked to virtualize an application such as SharePoint and you are able to lend the knowledge obtained in this book.

In general, IT managers and architects interested in virtualizing tier-1 applications in VMware vSphere environment should read this book. It will provide insights into the virtualization

technology and understanding of considerations and strategies for running mission-critical applications using VMware vSphere.

Although each of the major Microsoft tier-1 applications is covered in this book, these lessons are focused on virtualization and in no way can touch the level of detail that a dedicated book on the specific application can cover. Our assumption is that you have thorough knowledge of the Microsoft server applications that you plan to deploy in a virtual environment. We will not discuss application-specific design and architecture in this book. Several books are dedicated to each server application on these topics; we suggest perusing the Sybex website. The objective of this book is to discuss specific considerations when it comes to virtualizing these applications using VMware vSphere. As a result, you will also notice that we do not cover every detailed step when configuring these products in our proof-of-concept deployments.

The information shared in this book is based on the authors' personal experiences designing, implementing, testing, and virtualizing Microsoft Tier-1 applications. The authors have spent considerable time virtualizing these Microsoft Tier-1 applications using VMware vSphere 4. The information shared in this book in no way reflects VMware's official position and recommendations for virtualizing Microsoft Tier-1 applications and should not be construed as such.

What You Will Learn

In this book we show you the basic considerations for virtualizing Microsoft tier-1 applications. You will learn how to deploy these applications in VMware vSphere environment and monitor them on an ongoing basis. We will also provide details on other VMware solutions that simplify other virtualization tasks, such as virtual machine provisioning, patching, and backup. You will learn how to make your journey in the virtualization realm a much simpler one through the use of virtualization and the vSphere product feature set.

This book is not designed to teach you how to deploy these applications, and the proof of concepts in each chapter show one way to deploy them. It should be understood that there are a countless number of ways to design and deploy these applications.

With the proof of concepts, we try to give you a simple business case for designing and deploying the application. You should always evaluate your business and technical requirements for performance, operation, business continuity, and disaster recovery if required. This book shows you the vSphere products that are available for use in each of the aforementioned areas for design and deployment of these applications on the vSphere virtualization platform.

What You Need

Trying to describe virtualization is not an easy task. Trying to describe how to execute specific tasks in a virtualized environment without visual aids is near impossible. To gain the most out of this book, you will want to follow along and use the proof of concepts as guides to setting up and configuring your own virtualized Microsoft environment. VMware offers ESXi as a free download. If you don't have the required hardware for running ESXi, you may try VMware Server (another free download), VMware Workstation, or VMware Fusion (for the Mac lovers out there). The latter two will give you a month or so of free use on which you can build Virtual Machines and get hands-on experience with what is discussed in this book. Additionally, the Microsoft products used throughout this book can all be downloaded from Microsoft's website and used for an evaluation period.

Aside from a virtualization platform on which to begin working, an understanding of the tier-1 Microsoft applications covered in this book will be helpful. The applications covered each

have books in the 1,000-page ballpark because of the vast number of features available (look for the Mastering series published by Sybex). Covering each application end-to-end is not the intent here; rather, we cover what to look out for when you toss vSphere into the mix. As mentioned in the book, it is recommended that you not use the VMware Server and Workstation products for any deployments of these applications in a production environment. To take advantage of all the features covered in this book, the vSphere products we recommend include vSphere 4.0 and higher for your proof of concepts and virtual server deployments.

What Is Covered in This Book

Virtualizing Microsoft Tier 1 Applications with VMware vSphere 4 is organized as follows:

Chapter 1: VMware vSphere 4 Overview This chapter provides an overview of vSphere and details about its key features. It specifically points out features that are new in vSphere 4.

Chapter 2: Virtualizing Microsoft Server Applications This chapter discusses the topics that are common across these Windows server technologies and applications when virtualizing with VMware vSphere, such as virtualization technologies, hardware issues, virtual deployment strategies, licensing, data backups, patches, and updates.

Chapter 3: Virtualizing Windows Server 2008 This chapter takes a look at the requirements, capabilities, and virtual hardware limits to consider when building Windows Server 2008 virtual machines. You will learn how to quickly deploy consistent virtual machine images using the templates and tools found in vSphere.

Chapter 4: Virtualizing Active Directory This chapter discusses considerations and best practices for virtualizing Active Directory on the vSphere platform.

Chapter 5: Virtualizing Internet Information Services This chapter provides an understanding of some of the differences in the latest version of IIS and how new features can be leveraged when deployed in vSphere.

Chapter 6: Virtualizing Exchange Server This chapter takes a deep dive into the performance characteristics of Exchange 2007 and 2010 and shows how dramatic improvements to Exchange have made virtualizing the most widely used messaging system in use as an accepted and proven solution.

Chapter 7: Virtualizing SQL Server This chapter helps SQL Server DBAs and IT administrators who want to virtualize SQL Server using VMware vSphere. It will address their frequently asked questions and provide a proof-of-concept methodology for SQL Server virtualization with VMware vSphere; it also covers the considerations for virtualizing SQL Server on the vSphere platform and how to characterize existing SQL Server deployments for virtualization.

Chapter 8: Virtualizing Microsoft Office SharePoint Server 2007 This chapter breaks down the multitiered application and helps you fully understand how each piece fits into the vSphere infrastructure.

Chapter 9: Virtualizing Microsoft Remote Desktop Services This chapter reviews the Remote Desktop Services (RDS) roles and enhancements in Windows Server 2008 R2, covers key considerations for virtualizing RDS roles, discusses how to check application compatibility for RDS deployment, and discusses how to characterize and size RDS environments. We focus on the traditional session-based remote desktops provided by RDS in this chapter. The Virtual Desktop Infrastructure (VDI) enabled by RDS is beyond the scope of this book.

How to Contact the Authors

We welcome feedback from you about this book or about books you'd like to see from us in the future.

Charles A. Windom Sr. contributes to various communities, such as VMware's Exchange, Domino, and RIM community as well as Microsoft's Exchange Admin and General Discussion forums. Charles also occasionally blogs on TechPitch as well as other VMware blog sites. Questions for Charles can be directed to `cwindomsr@hotmail.com`.

Hemant Gaidhani authors his personal blog and virtualization-management blog on VMware Communities website. He also contributes to various VMware communities. He can be reached at `hemantgaidhani@yahoo.com`.

You can find Alex Fontana contributing to various communities, such as VMware's Exchange, Domino, and RIM community as well as Microsoft's Exchange Admin and General Discussion forums. Alex will also occasionally pop up on the ActiveDir.org and Sunbelt software Exchange mailing lists. Any questions for Alex can be directed to `alex_fontana@hotmail.com`.

Sybex strives to keep you supplied with the latest tools and information you need for your work. Please check its website at `www.sybex.com`, where we'll post additional content and updates that supplement this book if the need arises. Enter ***virtualizing microsoft tier 1*** in the Search box (or type the book's ISBN—**9780470563601**) to find the book's web page.

Chapter 1

VMware vSphere 4 Overview

VMware vSphere 4 is the new major version of VMware's flagship virtualization platform, VMware Infrastructure 3. A bare-metal hypervisor that enables full virtualization of industry-standard x86 hardware forms the foundation of this virtualization platform. In addition to this hypervisor, vSphere includes several advanced features that support innovative applications of virtualization technology. These features range from resource pooling, dynamic balancing of workloads, high availability, and disaster recovery.

VMware classifies all these vSphere features into this set of services:

- Infrastructure services
- Application services
- Management services

The infrastructure and application services are part of vSphere, and the management services are provided by VMware vCenter Server.

In this chapter, we will describe these services and provide details of their features. We will specifically call out features that are new in vSphere 4.

Infrastructure Services

vSphere infrastructure services are the core set of services that allows you to virtualize x86 servers (see Figure 1.1). First, these services abstract the physical x86 hardware resources, such as CPU, memory, storage, and network adapters, into virtual hardware to create virtual machines (VMs). Next, these services enable vSphere to transform resources from individual x86 servers into a shared computing platform with several operating systems and applications running simultaneously in different virtual machines. Finally, the infrastructure services provide several sophisticated features to optimize resources in such a shared environment. Figure 1.1 provides an overview of all the services in vSphere 4.

VMware vSphere provides the following types of infrastructure services:

- VMware vCompute
- VMware vStorage
- VMware vNetwork

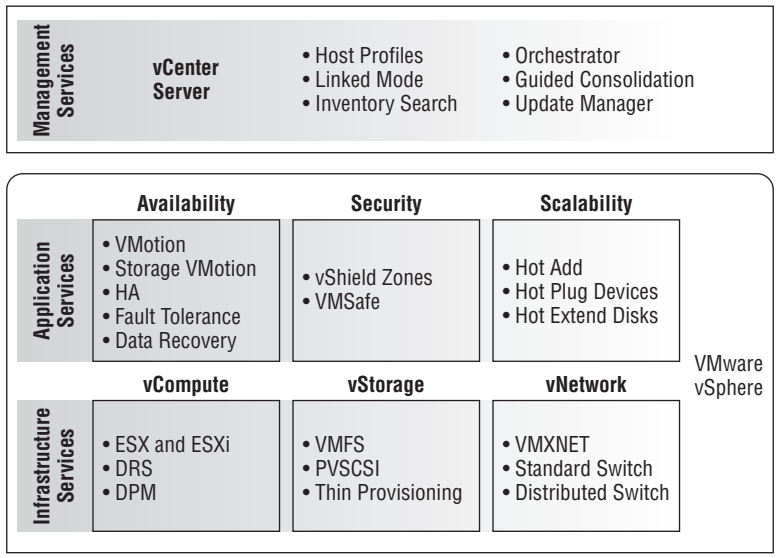

FIGURE 1.1 VMware vSphere overview

In the following sections, we'll provide a closer look at each type of infrastructure service.

VMware vCompute

VMware vCompute services virtualize CPU and memory resources in an x86 server. The vCompute services also aggregate these resources from several discrete servers into shared logical pools that can be allocated to applications running inside virtual machines.

The vCompute services comprise the following:

VMware ESX (and VMware ESXi) A bare-metal hypervisor that runs directly on server hardware. It supports different x86 virtualization technologies such as VMware-invented binary translation, hardware-assisted virtualization, and paravirtualization. VMware ESXi is a free version of ESX with a smaller footprint that minimizes the surface area for potential security attacks, making it more secure and reliable. ESX also includes several advanced CPU scheduling capabilities, as well as unique memory management features such as transparent page sharing and memory ballooning. These sophisticated features enable ESX to achieve higher consolidation ratios compared to its competition.

VMware Distributed Resource Scheduler (DRS) Extends the resource management features in ESX across multiple physical servers. It aggregates CPU and memory resources across many physical servers into a shared cluster and then dynamically allocates these cluster resources to virtual machines based on a set of configurable options. DRS makes sure that resource utilization is continuously balanced across different servers in the shared cluster.

VMware Distributed Power Management (DPM) Included with VMware DRS, DPM automates energy efficiency in VMware DRS clusters. It continuously optimizes server power consumption within each cluster by powering on or off vSphere servers as needed.

In the next sections, we will discuss each of these vCompute services in detail.

VMware ESX and ESXi

VMware ESX and ESXi are the most widely deployed virtualization hypervisors, and they form the robust foundation of VMware vSphere. VMware ESX and ESXi use bare-metal architecture; in other words, they install directly on the server hardware, without the need for a host operating system.

Virtualization Architectures

Virtualization products for x86 servers commonly use two types of architectures: a hosted architecture or a hypervisor architecture. The hosted, or *type 2*, virtualization products run on top of the Windows or Linux host operating system. The host operating system controls the access to the physical resources, and virtual machines run as applications alongside other software on the host machine. The VMware Workstation, Fusion, and Server products are examples of hosted virtualization architecture.

Bare-metal hypervisor, or *type 1*, virtualization products run directly on top of the hardware with direct access and control of the hardware's resources. Since they have direct access to the hardware resources rather than going through a host operating system, the hypervisor products are more efficient than hosted virtualization products, and deliver greater scalability, robustness, and performance. VMware ESX and ESXi are examples of bare-metal hypervisor architecture.

Virtualization Technologies

VMware ESX and ESXi offer a choice of three virtualization technologies (Figure 1.2):

- Binary translation
- Hardware-assisted virtualization
- Paravirtualization

Binary translation is the virtualization technique that VMware invented for x86 servers. The x86 processors were not designed with virtualization in mind. These processors have 17 CPU instructions that require special privileges and can result in operating system instability when virtualized. The binary translation technique translates these privileged instructions into equivalent safe instructions, thus enabling virtualization for x86 servers. Binary translation does not require any specific features in the x86 processors and hence enables you to virtualize any x86 server in the data center without modifying guest operating system and applications running on it.

Hardware-assisted virtualization relies on the CPU instruction set and memory management virtualization features that both AMD and Intel have recently introduced in the x86 processors. The first generation of these hardware-assisted virtualization processors, called *AMD-SVM* and *Intel-VT*, only supported CPU instruction set virtualization in the processors. This alone did not perform fast enough for all different workloads, compared to the binary translation technology.

Recently, AMD and Intel have introduced newer processors that also support memory management virtualization. You'll learn more about this in the next section. Virtualization using these second-generation hardware-assisted processors usually performs better than binary translation. Consequently, with the release of vSphere, VMware ESX, and ESXi now default to hardware-assisted virtualization out of the box, but you do have the choice to override this setting.

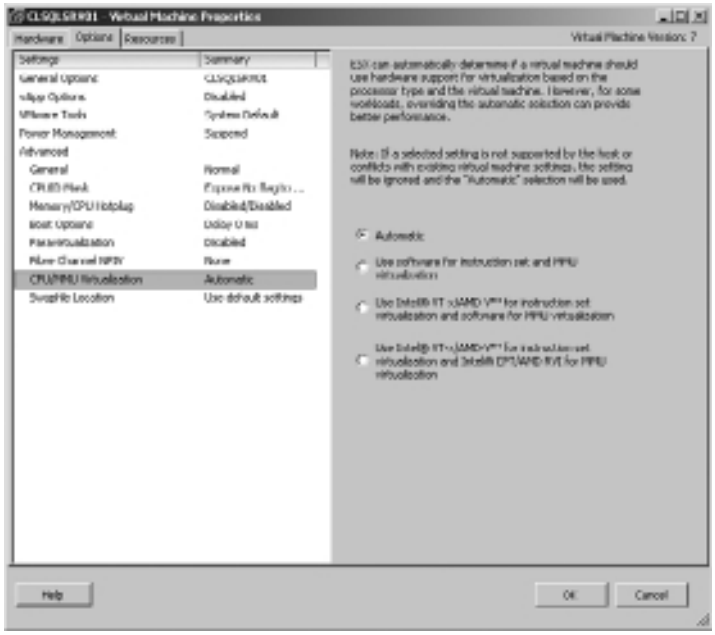

FIGURE 1.2
VMware vSphere virtualization technology options

VMware vSphere also supports *paravirtualized Linux guest operating systems*—Linux kernels that include Virtual Machine Interface (VMI) support—that are virtualization-aware. Because the VMI standard is supported out of the box in newer Linux kernels, there is no need to maintain separate distributions of Linux specifically for virtualization.

Hardware-Assisted Memory Virtualization

Memory management in virtual machines is challenging compared to physical machines, especially when it comes to virtual memory address translation. In a physical machine, the operating system uses page tables to translate memory addresses from an application's "virtual" space into the machine's physical memory addresses. Similarly, in a virtual machine, guest virtual memory addresses are translated to guest physical addresses using the guest OS's page tables. However, the guest OS does not have access to the physical machine memory; ESX controls the access to the actual physical memory. ESX performs the final translation to machine physical memory addresses by implementing a set of shadow page tables for each virtual machine. Creating/maintaining the shadow page tables adds both CPU and memory overhead. This overhead can be significant for virtual machines running several processes or using multiple virtual CPUs.

Both AMD and Intel have introduced hardware-assisted memory management capabilities to alleviate this situation. Processors supporting hardware-assisted memory management implement an additional level of page tables in hardware. These hardware page tables keep track of guest

physical to machine memory address translations, which used to be maintained inside shadow page tables within ESX. Offloading this memory management to hardware has two benefits: hardware page table processing is faster than software implementation, and ESX can use the freed CPU cycles for more workload-related processing.

AMD calls its hardware-assisted memory management feature *rapid virtualization indexing* (RVI), while Intel terms its implementation *extended page tables* (EPT). ESX has supported AMD RVI since version 3.5. The support for Intel EPT was introduced in ESX 4.0.

The performance benefits of hardware-assisted memory management are achievable only if page table entries are located in hardware page tables. Remember that the real estate on a processor chip is at a premium and hence limits the size of hardware page tables. If a page table entry is not found in the hardware page tables, the associated translation lookaside buffer (TLB) miss can result in more expensive processing compared to software shadow page tables implemented by ESX. You can reduce the number of TLB misses by using large memory pages. ESX has been supporting large memory pages since version 3.5. Together, hardware-assisted memory management and large memory pages will provide better performance.

Processor Scheduling

VMware vSphere includes a sophisticated CPU scheduler that enables it to efficiently run several machines on a single ESX host. The CPU scheduler allows you to over-commit available physical CPU resources; in other words, the total number of virtual CPUs allocated across all virtual machines on a vSphere host can be more than the number of physical CPU cores available. The virtual machines are scheduled on all available physical CPUs in a vSphere host by default or can be affinitized or pinned to specific physical CPUs. The ESX CPU scheduler will also guarantee that a virtual machine only uses CPU cycles up to its configured values. When scheduling virtual CPUs allocated to virtual machines, the CPU scheduler uses a proportional-share scheduling algorithm that also takes into account user-provided resource specifications such as shares, reservations, and limits. Maintaining CPU resource allocation fairness among a number of virtual machines running on a vSphere host is a key aspect of ESX processor scheduling.

Starting with the Virtual Infrastructure 3 (VI3) release, ESX has gradually shifted from "strict" to "relaxed" co-scheduling of virtual CPUs. Strict co-scheduling required that a virtual machine would run only if all its virtual CPUs could be scheduled to run together. With relaxed co-scheduling, ESX can schedule a subset of virtual machine CPUs as needed without causing any guest operating system instability.

The ESX CPU scheduler is also aware of different processor topology architectures such as nonuniform memory access architecture (NUMA) nodes and hyperthreading.

The ESX 4.0 scheduler further improves on these capabilities by adding the following enhancements:

- More optimizations to relaxed co-scheduling of virtual CPUs, especially for SMP VMs (virtual machines with multiple virtual CPUs)
- New finer-grained locking to reduce scheduling overheads in cases where frequent scheduling decisions are needed
- Processor cache topology awareness and optimizations to account for newer processor cache architectures
- Improvements in interrupt delivery efficiency and the associated processing costs

Advanced Memory Management

VMware vSphere uses several advanced memory management features to efficiently use the physical memory available. These features make sure that in a highly consolidated environment virtual machines are allocated the required memory as needed without impacting the performance of other virtual machines. These advanced features include the following:

Memory over-commitment Similar to CPU over-commitment, memory over-commitment improves memory utilization by enabling you to configure virtual machine memory that exceeds the physical server memory. For example, the total amount of memory allocated for all virtual machines running on a vSphere host can be more than the total physical memory available on the host.

Transparent page sharing Transparent page sharing uses available physical memory more efficiently by sharing identical memory pages across multiple virtual machines on a vSphere host. For example, multiple virtual machines running Windows Server 2008 will have many identical memory pages. ESX will store a single copy of these identical memory pages in memory and create additional copies only if a memory page changes.

Memory ballooning Memory ballooning dynamically transfers memory from idle virtual machines to active ones. It puts artificial memory pressure on idle virtual machines, forcing them to use their own paging areas and release memory. This allows active virtual machines in need of memory to use this memory. Keep in mind that ESX will ensure that a virtual machine memory usage cannot exceed its configured memory.

Large memory pages Newer x86 processors support the use of large 2 MB memory pages in addition to the small 4 KB pages. Operating systems rely on the translation lookaside buffers inside the processor to translate virtual to physical memory addresses. Larger page sizes mean that a TLB cache of the same size can keep track of larger amounts of memory, thus avoiding the costly TLB misses. Enterprise applications such as database servers and Java virtual machines commonly use large memory pages to increase TLB access efficiency and improve performance. ESX supports the use of large memory pages in virtual machines and backs up with its own large memory pages to maintain efficient memory access.

Resource Management

VMware vSphere allows you to establish minimum, maximum, and proportional resource shares for CPU, memory, disk, and network bandwidth for virtual machines. The minimum resource setting or reservation guarantees the amount of CPU and memory resources for a virtual machine, while the maximum resource setting or limit caps the amount of CPU and memory resources a virtual machine can use. The proportional resource allocation mechanism provides three levels—normal, low, and high—out of the box. These settings help configure virtual machine priority for CPU and memory resources relative to each other. These can be set at the resource pool level and are inherited or overridden at the individual virtual machine level. You can leverage these resource allocation policies to improve service levels for your software applications. The key advantage of these settings is that you can change resource allocations while virtual machines are running, and the changes will take place immediately without any need to reboot.

You need to be careful when assigning the minimum settings or reservations because they guarantee resources to a virtual machine. If too much CPU and memory resources are reserved, you may not be able to start virtual machines.

New Virtual Hardware Generation

In vSphere 4.0, VMware has upgraded the virtual hardware from version 4 to version 7. This generation of virtual hardware adds the following features:

Serial attached SCSI (SAS) virtual device for Microsoft Cluster Service This virtual device is needed to support running Windows Server 2008 in a Microsoft Cluster Service configuration. Later chapters in this book cover setting up a Microsoft Cluster Service configuration using Windows Server 2008 and will demonstrate the use of this device.

IDE virtual device This virtual device is recommended for older operating systems that do not support SCSI drivers.

VMXNET generation 3 VMware introduces the third generation of their paravirtualized virtual networking adapter in vSphere 4. Refer to the "VMware vNetwork" section later in this chapter for details.

Virtual machine hot-plug support The new virtual hardware generation enables you to hot plug virtual devices to a virtual machine without having to power it off. You can hot add and remove virtual CPUs, hot add and remove network cards and disks, and hot add memory to a virtual machine when using virtual hardware version 7. The support for the CPU and memory hot-add plug-in depends upon the guest operating system support.

New virtual machines created in the vSphere environment use virtual hardware version 7 by default. However, vSphere can run virtual machines created on hosts running ESX Server versions 2.*x* and 3.*x*.

You can find the virtual hardware version either on the virtual machine Summary tab or at the top left of the Virtual Machine Properties window. To convert your existing virtual machines, you first need to update VMware Tools to the latest version. You can then upgrade the virtual hardware used by the virtual machine. Right-click the virtual machine, and the context menu should provide you with an option to perform this upgrade (Figure 1.3).

FIGURE 1.3
VM hardware upgrade

NOTE Virtual machines using virtual hardware version 7 features are not compatible with ESX/ESXi releases prior to version 4.0. The virtual machine upgrade process is irreversible, and you will see a warning window during the upgrade steps, as shown here.

Distributed Resource Scheduler

VMware DRS allows you to manage physical resources distributed across multiple ESX servers. Using DRS, you can aggregate CPU and memory resources from up to 32 ESX servers to create a shared pool of resources, appropriately called *resource pools*. You can then organize these resource pools to create a flexible hierarchy to reflect business priorities. DRS also allows you to extend the resource management capabilities of a single ESX server such as shares, reservations, or limits to all virtual machines within these resource pools (Figure 1.4). For example, you can assign higher shares of the total resources to the production resource pool compared to a test and development resource pool. Likewise, you can guarantee fixed CPU and memory (reservations) for business-critical applications within that production resource pool.

FIGURE 1.4
VMware vSphere resource pools

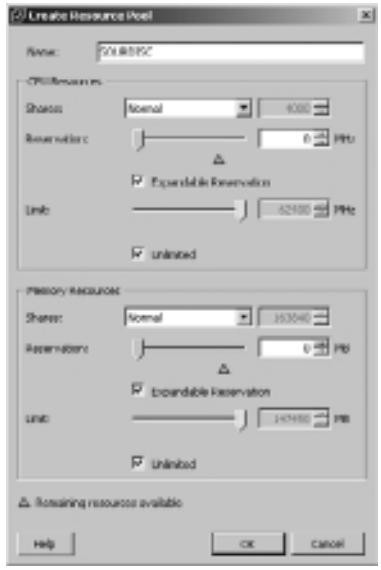

DRS also allows you to define rules and policies for virtual machines' resource allocations (Figure 1.5). For example, you can define an affinity rule to make sure that all virtual machines in a given application stack always run on the same server. All network communication for such co-located virtual machines takes place in memory and can benefit application performance. Alternatively, you can define an anti-affinity rule to ensure that specific virtual machines always

run on different servers. You can use this to avoid a single point of failure and increase availability for application components such as web servers in a load-balanced farm.

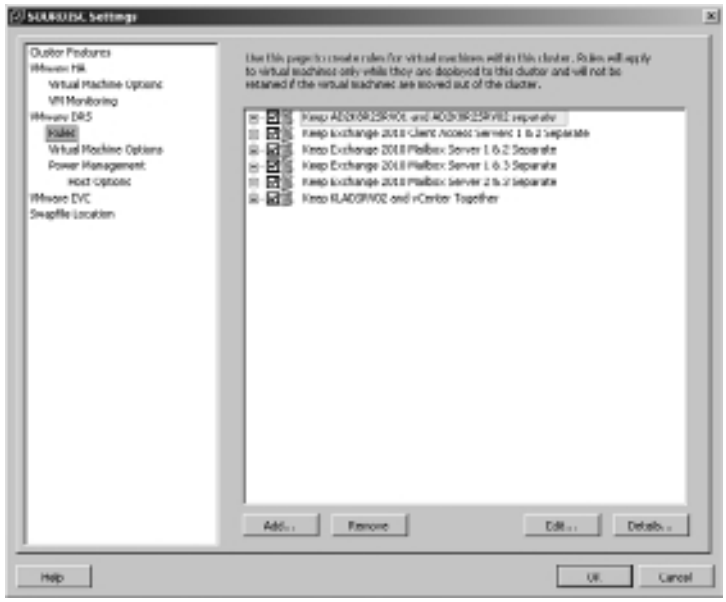

FIGURE 1.5
VMware vSphere DRS rules

VMware DRS will help you to load balance the resource utilization across various ESX servers within resource pools. It continuously monitors utilization across resource pools and then generates recommendations to optimally place virtual machines within these resource pools. While doing this, DRS also takes into account any priority settings, such as shares, reservations, and limits, as well as resource allocation rules and policies you may have defined. You can configure to execute these virtual machine placement recommendations automatically or manually (Figure 1.6). VMware DRS does this smart resource allocation during the initial virtual machine placement and on continuous basis in real time while the virtual machine is running:

Initial placement When a virtual machine is first powered on, VMware DRS either places the virtual machine on the most appropriate ESX server automatically or makes a recommendation for you to act on.

Continuous optimization The virtual machine resource requirements change over time, and VMware DRS dynamically responds to these changing needs by migrating virtual machines between ESX servers using VMware VMotion without disrupting end users. Alternatively, you can configure DRS to use manual mode so it only makes recommendations that you can then choose to act on.

By leveraging VMotion, VMware DRS also simplifies planned maintenance on physical servers without disrupting virtual machines and end users. When you place a physical server in maintenance mode, VMware DRS identifies alternative servers where the virtual machines can run. Based on the automation mode settings, either the virtual machines are automatically moved to use the alternative servers or the system administrator performs the move manually using the VMware DRS recommendations as a guideline.

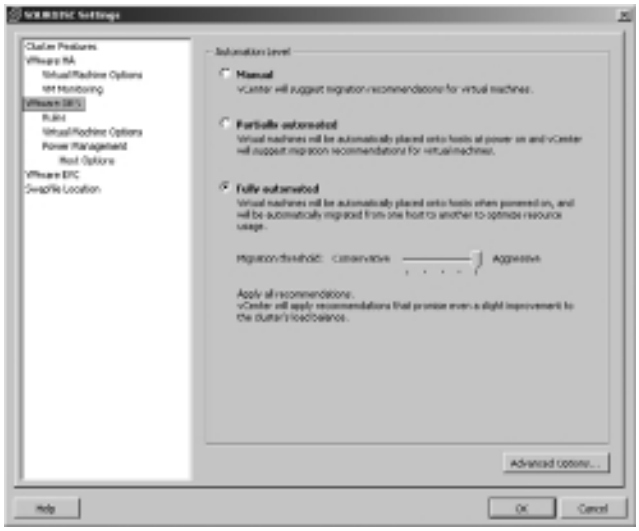
Cluster Features
VMware HA
Virtual Machine Options
VM Monitoring
VMware DRS
Rules
Virtual Machine Options
Power Management
Host Options
VMware EVC
Swapfile Location
Automation Level
Manual
vCenter will suggest migration recommendations for virtual machines.
Partially automated
Virtual machines will be automatically placed onto hosts at power on and vCenter will suggest migration recommendations for virtual machines.
Fully automated
Virtual machines will be automatically placed onto hosts when powered on, and will be automatically migrated from one host to another to optimize resource usage.
Migration threshold: Conservative Aggressive
Apply all recommendations.
vCenter will apply recommendations that promise even a slight improvement to the cluster's load balance.
Advanced options...
Help
OK
Cancel

FIGURE 1.6
VMware vSphere DRS automation levels

DISTRIBUTED POWER MANAGEMENT

VMware DPM enables you to reduce energy consumption in the data center by optimizing the placement of virtual machines within a cluster. When resource utilization in a DRS cluster goes down, DPM consolidates virtual machines within the cluster on fewer ESX servers and powers off unused servers to conserve energy. When resource requirements increase, DPM brings ESX servers back online to ensure that virtual machines get appropriate resources and service levels can be maintained (Figure 1.7).

DPM uses technologies such as IPMI, iLO, and Wake on LAN to control server power states. ESX/ESXi 4.0 also supports the Enhanced Intel SpeedStep and Enhanced AMD PowerNow! CPU power management technologies. These technologies leverage dynamic voltage and frequency scaling (DVFS) to reduce power consumption.

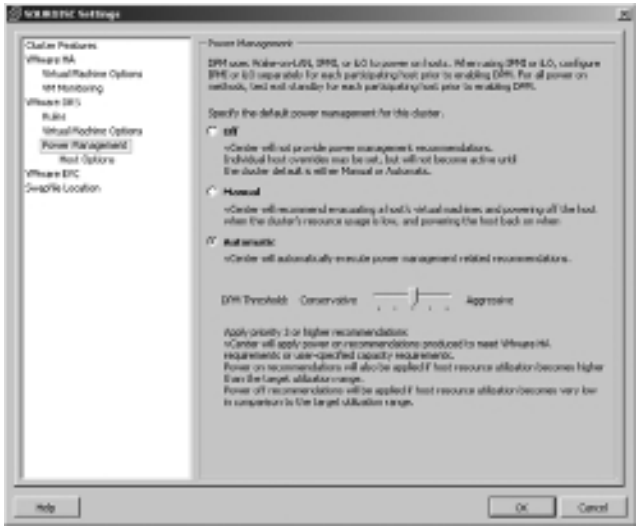

FIGURE 1.7
VMware vSphere DPM power management

VMware vStorage

VMware vStorage includes a number of advanced features to efficiently manage and use storage in virtualized environments while trying to hide the complexities of the underlying storage systems. The key features include the following:

- VMware vStorage *Virtual Machine File System* (VMFS) is a custom-designed cluster file system specifically optimized to store and manage virtual machines. It allows the efficient sharing of back-end storage by multiple ESX servers and is the key enabler for VMware features such as VMotion, Storage VMotion, DRS, VMware High Availability, and Fault Tolerance.
- VMware vStorage *thin provisioning* enables you to over-commit storage capacity similar to CPU and RAM over-commitment. Thin provisioning initially allocates only as much storage capacity as needed and grows as more data is stored in the virtual disk. This increases storage utilization allowing you to defer storage purchases until they are really required. This can significantly reduce an organization's storage budget.
- vSphere introduces a high-performance, *paravirtualized SCSI (PVSCSI) storage adapter* that offers greater throughput and lower CPU utilization for virtual machines. This is best suited for environments that run very I/O-intensive guest applications.
- VMware vStorage offers APIs for data protection, multipathing, and storage replication adapters so that storage partners can integrate their solutions with vSphere.

In the next sections, we will discuss each of these vStorage services in detail.

Virtual Machine File System

VMware created a cluster file system, VMFS, specially designed to store and manage virtual machines. VMFS is optimized to support large files associated with virtual disks, thus enabling encapsulation of an entire virtual machine in a set of files. Using VMFS, you can place these virtual machine files on a shared storage and allow multiple ESX servers to concurrently read and write to this shared storage.

By managing concurrent access to the shared back-end storage, VMFS enables the foundation for key VMware features such as VMotion, Storage VMotion, DRS, VMware High Availability, and Fault Tolerance. As virtual machines are migrated to or restarted on different ESX servers, VMFS ensures that individual ESX servers are not single points of failure and helps DRS to balance resource utilization across multiple servers. VMFS uses on-disk file locking to ensure that the same virtual machine is not powered on by multiple servers at the same time.

VMFS also acts a logical volume manager by providing an interface to different types of storage such as Fibre Channel SAN, iSCSI SAN, and NAS. VMFS hides the complexities of underlying storage systems and, irrespective of the storage type, simplifies storage management using automatic discovery and mapping of LUNs to a VMFS volume. You can connect or disconnect a VMware ESX server from a VMFS volume without impacting other VMware ESX hosts. vSphere also adds dynamic growth capabilities to VMFS without the need for any downtime. These new capabilities include hot expansion of VMFS volumes and virtual disks stored in VMFS.

vStorage VMFS Volume Grow The VMFS Volume Grow capability in vSphere allows you to dynamically expand the size of an existing data store that resides on a VMFS volume without disrupting running virtual machines. It complements the dynamic LUN expansion capability that exists in many storage array offerings today. After you expand the LUN where a data store resides through an array management utility, you can use VMFS Volume Grow to expand the VMFS extent on the expanded LUN. You can verify the increased VMFS volume (data store) size from vCenter Server.

For earlier versions of ESX, you have to use VMFS spanning across multiple LUNs to increase the size of an existing VMFS volume. First, you expand the LUN upon which the VMFS volume resides; next, you create a separate disk partition in that additional storage space and add the new partition as if you were adding a second LUN to the VMFS volume.

Hot extend for virtual disks Hot extend for virtual disks allows you to add virtual storage to running virtual machines without any downtime. You can use hot extend for VMFS flat virtual disks using persistent mode and for ones that do not have any VMFS snapshots. You will need to run guest operating system tools for it to start using the additional storage. Together with the VMFS Volume Grow capability, this feature provides a very flexible and dynamic way to manage storage capacity growth.

Virtual Disk Thin Provisioning

VMware virtual disk thin provisioning enables over-commitment of storage capacity similar to CPU and RAM over-commitment. It allows the storage administrator to configure more virtual machine storage capacity than the actual physical storage currently available. This is possible because thin provisioning enables virtual machines to utilize storage space on an as-needed basis. When a virtual disk is initially allocated, it is assigned 1 MB of storage space in the data store. As that disk space is used up, additional 1 MB chunks of storage are allocated for the virtual disk so that the underlying storage demand will grow as its size increases. This dynamic allocation reduces storage over-provisioning and increases storage utilization, allowing you to defer storage purchases until they are really required. This can significantly reduce an organization's storage budget.

VMware's desktop products such as VMware Workstation have long provided the feature to allow virtual machines to allocate storage space as needed. In Virtual Infrastructure 3, thin provisioning was used by default for virtual disks created on NFS data stores and was available for block-based data stores through the command line. With VMware vSphere, vCenter now fully supports virtual disk thin provisioning for all virtual disks when deploying or migrating virtual machines (Figure 1.8).

vCenter Server 4.0 continuously monitors the storage allocations and tracks it against storage capacity so that it can generate alerts and alarms to warn vSphere administrators against any future "out of space" situations.

NOTE Virtual disk thin provisioning should not be confused with the same thin provisioning that an array vendor might offer. In fact, with vSphere, you now have the capability of doing thin provisioning at the data store level in addition to doing thin provisioning at the storage array.

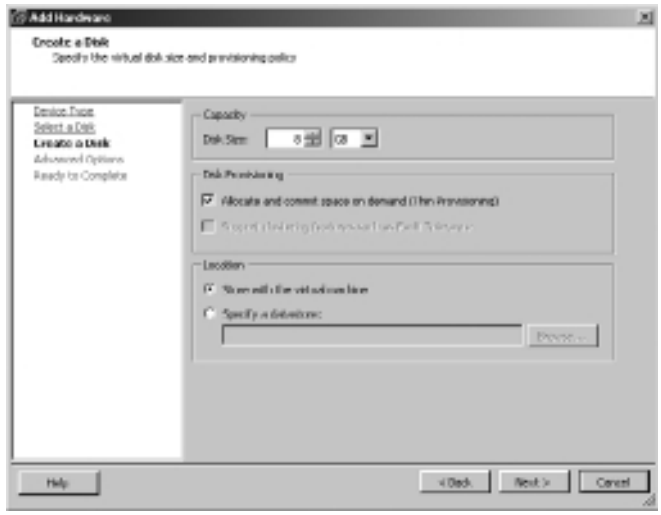

FIGURE 1.8
Creating a thin-provisioned virtual hard disk

VMware Paravirtualized SCSI

Early versions of ESX supported virtual SCSI adapters that emulated BusLogic and LSI Logic hardware storage adapters. This full virtualization provided broad compatibility with guest operating systems supported by ESX. However, this prevents ESX from taking advantage of performance optimizations that can be achieved with paravirtualized devices such as VMXNET network adapters. vSphere introduces a new virtual storage adapter called PVSCSI, as shown in Figure 1.9, which extends to the storage stack performance gains typically associated with other paravirtual devices. In that respect, the PVSCSI adapter is similar to the paravirtualized network adapter VMXNET that is available in ESX. As with other paravirtual devices, the PVSCSI adapter improves I/O efficiency by using optimizations such as batching the processing of I/O requests and I/O completion interrupts and reducing the cost of virtual interrupts. The PVSCI adapter also benefits from a reduced number of context switches between the guest operating system and ESX virtual machine monitor.

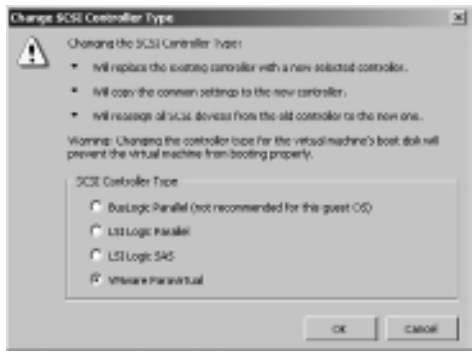
Change SCSI Controller Type
Changing the SCSI Controller Type:
Will replace the existing controller with a new selected controller.
Will copy the common settings to the new controller.
Will reassign all virtual devices from the old controller to the new one.
Warning: Changing the controller type for the virtual machine's boot disk will prevent the virtual machine from booting properly.
SCSI Controller Type:
BusLogic Parallel (not recommended for this guest OS)
LSI Logic Parallel
LSI Logic SAS
VMware Paravirtual
OK
Cancel

FIGURE 1.9
Paravirtualized SCSI controller

The performance benefits of the PVSCSI driver are visible for virtual machines issuing more than 2,000 I/O requests per second. For lower I/O throughput, VMware recommends that you continue to use the default or primary adapter. For example, LSI Logic is the default primary adapter for virtual machines with Microsoft Windows 2008 guest operating systems. You can use it for the virtual disk that hosts the system software (boot disk) and a separate PVSCSI adapter for the disk that stores user data, such as a database.

In vSphere 4.0, you cannot use the PVSCI adapter for a boot partition, but subsequent versions are expected to support this. The PVSCSI driver currently works with only certain guest OS versions such as Windows Server 2003, Windows Server 2008, and RHEL 5. It can also be shared by multiple virtual machines running on a single ESX, unlike the VMDirectPath I/O, which will dedicate a single adapter to a single virtual machine.

vStorage APIs

vSphere introduces several storage APIs to enable integration with leading storage hardware and software products for data protection, high availability, and disaster recovery. Storage partners have written plug-ins to vSphere using a pluggable storage architecture, which is an open modular framework to leverage vendor-specific capabilities for better performance. You can use these partner plug-ins for better flexibility and visibility when configuring storage resources for your deployment. The vStorage APIs include the following:

The vStorage APIs for Multipathing These provide an I/O multipathing framework for storage partners to create Multipathing Extensions Modules that plug in to VMware ESX/ESXi to deliver storage path failover and storage I/O throughput optimized for partners' storage arrays.

The vStorage API for Data Protection This enables backup tools to directly connect the ESX servers and the virtual machines running on them without any additional software installation. This API allows backup tools to do efficient incremental, differential, and full-image backups and restores of virtual machines. To avoid any service disruption, this API also makes it possible to offload backup processing from ESX servers.

Storage Performance and Scalability

VMware has made significant performance and scalability improvements to the storage stack in vSphere. These enhancements apply to all supported storage protocols: Fibre Channel SAN, iSCSI SAN, and NFS. Together with the new paravirtualized SCSI driver, these storage stack optimizations dramatically improve storage I/O performance, in terms of both I/O throughput and the CPU cost per I/O. These are the key storage performance improvements:

Improved I/O efficiency With VMware vSphere 4, like earlier versions of ESX, you can achieve I/O throughput levels that are limited only by the capabilities of the underlying storage system and the storage connectivity link speed. Because of storage stack optimizations, vSphere 4 uses fewer CPU cycles to achieve these throughput levels.

Software iSCSI and NFS support with jumbo frames Using jumbo frames is a recommended best practice to improve performance for Ethernet-based storage. Earlier ESX versions supported jumbo frames only for networking. In vSphere, you can now leverage jumbo frames for both NFS and iSCSI storage whether you use 1 Gbps or 10 Gbps NICs.

iSCSI support improvements In vSphere 4, VMware has rewritten the entire iSCSI software initiator stack for both software iSCSI (that is, in which the iSCSI initiator runs at the ESX layer) and hardware iSCSI (that is, in which ESX leverages a hardware-optimized iSCSI HBA). As a result, both software and hardware iSCSI in vSphere 4 provide better throughput and CPU efficiency when compared to the earlier 3.5 version.

VMDirectPath I/O for Storage

VMDirectPath I/O is a new capability in vSphere that enables virtual machines to directly access the underlying physical I/O devices. When using VMDirectPath I/O for storage, there is a one-to-one mapping between an HBA and a VM, and you are not allowed to share an HBA by more than one VM. VMDirectPath is designed to handle the I/O needs of special-purpose I/O appliances and I/O-intensive virtual machines. By accessing the I/O devices directly and bypassing the hypervisor, the guest OS enhances CPU efficiency in handling the high I/O workloads. However, this I/O throughput scalability comes at the cost of other virtualization features. Features such as VMotion, hardware independence, and sharing of physical I/O devices are not available for virtual machines using VMDirectPath I/O.

In vSphere 4.0, VMDirectPath I/O is experimentally supported for the following storage I/O devices:

- QLogic QLA25xx 8 Gb Fibre Channel adapters
- Emulex LPe12000 8 Gb Fibre Channel adapters
- LSI 3442e-R and 3801e (1068 chip based) 3 Gb SAS adapters

VMware vNetwork

VMware vNetwork provides features to help you deploy and manage enterprise-class virtual networking that can communicate with external physical networks. It comprises the following key features:

- *Virtual network adapters* enable network virtualization by allowing virtual machines to network like physical machines do. VMware provides three types of virtual network adapters that virtual machines can use. vSphere introduces the third generation of the paravirtualized in-guest network drivers (VMXNET3) for enhanced network I/O performance.
- The VMware vNetwork *Standard Switch* enables you to create a virtual network between virtual machines within a single VMware ESX/ESXi host as well as those on the outside physical network. These virtual switches support the same networking protocols as physical switches and implement enterprise-class features such as VLANs and hardware NIC teaming for availability and performance.
- The VMware vNetwork *Distributed Switch* moves beyond per-host network configuration and simplifies networking management across multiple hosts in VMware vSphere environments from one central interface. It also enables third-party distributed virtual switches such as the Cisco Nexus 1000V Series virtual switch to be used in VMware vSphere environments so that network administrators can use familiar interfaces when managing virtual networking.

We will discuss each of these vNetwork services in detail in the next sections.

Virtual Network Adapters

VMware supports three types of virtual network adapters for virtual machines. The AMD Lance PCNet32 and Intel E1000 physical network adapters provide full virtualization that is compatible with most guest operating systems. VMware also provides the high-performance paravirtualized network adapter VMXNET for networking virtual machines. In vSphere 4, VMware introduces VMXNET 3, the third generation of this paravirtualized network adapter (Figure 1.10). VMXNET3 builds upon the earlier versions, VMXNET and Enhanced VMXNET, by adding these features:

- MSI/MSI-X support (subject to guest operating system kernel support)
- Receive-side scaling (supported in Windows 2008 when explicitly enabled through the device's Advanced Configuration tab)
- IPv6 checksum and TCP segmentation offloading (TSO) over IPv6
- VLAN offloading
- Large TX/RX ring sizes (configured from within the virtual machine)

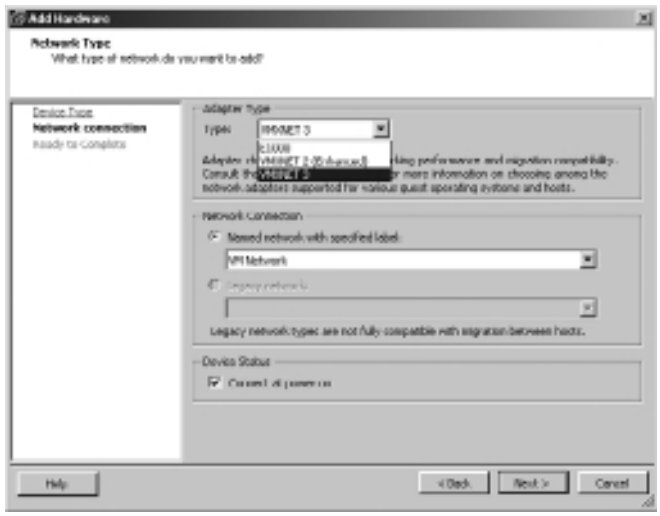

Figure 1.10 Setting a VMXNET3 network adapter

VMware includes two other network adapters, vswif and vmknic, for ESX/ESXi Service Console and VMkernel, respectively. All these virtual network adapters support both IPv4 and IPv6. The support for IPv6 for the ESX/ESXi VMkernel and Service Console was added in vSphere 4. vCenter Server 4 also supports IPv6 so that vSphere customers can manage mixed IPv4/IPv6 environments. IPv6 support for network storage is considered experimental in vSphere 4.0.

vNetwork Standard Switch

Virtual switches allow virtual machines on the same ESX Server host to communicate with each other using virtual ports and with an external network using uplink ports. These switches do

not need any additional networking hardware and are implemented in software based on the configuration you specify. Virtual switches are similar to physical switches in several ways:

- They use same networking protocols.
- They support VLANs compatible with standard VLAN implementations.
- They support Layer 2 forwarding.
- They support offloading features for TCP checksum and segmentation.
- They support Layer 2 security policies such as locking down MAC address changes.

The VMware virtual switch is also capable of binding multiple virtual NICs together, similar to NIC teaming in physical servers, to offer high availability and throughput for the virtual machines.

With more than 1,000 virtual ports per switch, you can support a large number of virtual machines per single virtual switch. Each virtual switch is isolated, and you cannot connect multiple virtual switches within the same vSphere host. This helps improve security for virtual networks, in addition to the Layer 2 security features listed earlier.

You can compare the features available in VMware virtual switches at `www.vmware.com/products/vnetwork-distributed-switch/features.html`.

vNetwork Distributed Switch

With vSphere, VMware has introduced the vNetwork Distributed Switch (vDS) that aggregates virtual switches from multiple vSphere hosts in a single distributed network switch. This eliminates the restriction of managing virtual networking for each vSphere host separately, and it simplifies network maintenance for the entire vSphere cluster. vDS provides a centralized interface from VMware vCenter Server for provisioning, administration, and monitoring of virtual networking for your entire data center. This can significantly reduce ongoing network maintenance activities and allow you to quickly scale up networking capacity. vDS also enables the following features:

- Network VMotion
- Bidirectional traffic shaping
- Third-party virtual switch support with the Cisco Nexus 1000V Series virtual switch

Network VMotion

Network VMotion maintains virtual machine networking state (for example, counters and port statistics) as the virtual machine moves from one host to another on a vDS. As a result, you have a consistent view for the virtual network interface regardless of which vSphere host a virtual machine is located on or how frequent a virtual machine is migrated by VMotion. This is very helpful in monitoring and troubleshooting any network-related activities in large-scale vSphere deployments.

Bidirectional Traffic Shaping

vNetwork Standard Switches allow you to set up traffic shaping for egress or transmit (from virtual machine to network) traffic. vDS expands this capability to include bidirectional traffic shaping. Both egress (from virtual machine to network) and ingress (from network into virtual machine) traffic-shaping policies can now be applied on DV port group definitions. You can use the following three characteristics to define traffic-shaping policies:

- Average bandwidth
- Peak bandwidth
- Burst size

You can leverage traffic shaping to limit the traffic to or from a virtual machine or group of virtual machines to protect either a virtual machine or other traffic in an over-subscribed network.

Third-Party Virtual Switch Support with the Cisco Nexus 1000V Series Virtual Switch

The vDS includes support for third-party distributed virtual switches. Cisco collaborated with VMware to leverage this extensibility to develop the Cisco Nexus 1000V Series virtual switch.

Both the Cisco Nexus 1000V Series virtual switch and the VMware vNetwork Distributed Switch use the same distributed switching model.

Virtual Ethernet modules (VEMs) Each ESX host implements VEMs, which are the switching data planes, and provide the frame-forwarding capabilities. These VEMs leverage the ESX host APIs and can support the same physical NICs and hardware compatibility list (HCL) as the VMware Standard Switch and vNetwork Distributed Switch.

Virtual supervisor modules (VSMs) The Cisco NX-OS operating system implements VSMs. They provide the control plane function for the VEMs and can exist as a guest VM or standalone appliance. VSMs allow you to use the familiar Cisco command-line interface (CLI) for management and configuration. You can also use vSphere Client to communicate with VSMs and vCenter Server for optional management and configuration.

The Cisco Nexus 1000V offers an expanded feature set compared to the VMware vNetwork Distributed Switch and is similar to that provided by the physical Cisco Catalyst and Nexus switches. You can find more information on the Cisco Nexus 1000V at `http://cisco.com/go/nexus1000v`.

Network Performance and Scalability

vSphere includes several performance enhancements to the network stack. These enhancements come in two ways: support for various offload technologies and optimizations to the existing network stack processing. It is a common performance optimization practice to offload various network processing tasks to reduce the CPU overhead associated with processing network I/O. The VMXNET3 network adapter supports performance offload technologies such as TCP Segmentation Offloading, TCP/IP Checksum Offload, and Large Receive Offload, as well as other optimizations like jumbo frames. vSphere 4 also includes optimizations to the network

stack, such as NetQueue, that can saturate even 10 Gbps links for both transmit and receive-side network I/O. You will also notice a significant increase in iSCSI throughput and maximum network throughput for VMotion because of the VMkernel TCP/IP stack optimizations in vSphere 4.

VMDirectPath I/O for Networking

VMDirectPath I/O for Networking is a new capability in vSphere that enables virtual machines to directly access underlying physical I/O devices. Similar to VMDirectPath I/O for Storage, VMDirectPath is designed for special-purpose I/O appliances and network I/O-intensive virtual machines that require very efficient network stack processing for high throughput but do not need to support additional virtualization features such as VMotion, Fault Tolerance, and suspend/resume.

In vSphere 4, VMDirectPath I/O for Networking is supported for the following devices:

- Intel 82598 10 Gigabit Ethernet Controller
- Broadcom 57710 and 57711 10 Gigabit Ethernet Controller

Application Services

Virtualization enables interesting use cases such as VMotion and VMware HA that can used by any application running inside virtual machines. VMware vSphere application services build upon such use cases to deliver enterprise readiness features for applications running inside VMware virtual machines. These services are expected to enhance the service levels for virtualized applications more easily compared to physical deployments.

VMware vSphere provides the following types of application services:

- Availability
- Security
- Scalability

We'll take a closer look at each type of application service in the next sections.

Availability

Improving availability for applications is probably the most innovative and exciting use of virtualization technology. With availability services in vSphere, you can lower both planned and unplanned downtime for all applications running inside VMware virtual machines. Furthermore, vSphere enables this high availability without the need for complex hardware or software clustering solutions.

To minimize service disruptions because of *planned* hardware downtime, VMware vSphere includes the following availability services:

VMware VMotion Using VMware VMotion, you can migrate running virtual machines from one vSphere server to another without impacting the applications running inside virtual machines. The end users do not experience any loss of service. You can leverage VMotion to

move virtual machines off a vSphere server for any scheduled hardware maintenance without the need for any application downtime.

VMware Storage VMotion VMware Storage VMotion enables similar functionality at the storage level. You can migrate virtual disks of running virtual machines from one storage array to another with no disruption or downtime. Storage VMotion will help you avoid any application downtime because of planned storage maintenance or during storage migrations.

To reduce service disruptions because of *unplanned* hardware downtime, VMware vSphere availability services include the following features:

VMware High Availability (HA) If a virtual machine goes down because of hardware or operating system failures, VMware HA automatically restarts the virtual machine on another ESX server within minutes of the event. VMware HA provides a much simpler and cost-effective high-availability solution compared to traditional clustering solutions.

VMware Fault Tolerance (FT) VMware FT improves high availability beyond VMware HA. By maintaining a shadow instance of a virtual machine and allowing immediate failover between the two instances, VMware FT avoids even the virtual machine reboot time required in the case of VMware HA. Thus, it prevents any data loss or downtime even if server hardware fails. Like VMware HA, it can be a cheaper and simpler alternative to traditional clustering solutions.

VMware Data Recovery VMware Data Recovery enables a simple disk-based backup and restore solution for all of your virtual machines. It does not require you to install any agents inside virtual machines and is completely integrated into VMware vCenter Server. VMware Data Recovery leverages data deduplication technology to avoid saving duplicate storage blocks twice, thus saving both backup time and disk space.

VMware VMotion

VMotion enables live migration of running virtual machines from one ESX server to another with no downtime (Figure 1.11). This allows you to perform hardware maintenance without any disruption of business operations. The migration of the virtual machine is quite seamless and transparent to the end user. When you initiate a VMotion, the current state of the virtual machine, along with its active memory, is quickly transferred from one ESX server to another over a dedicated network link, and the ESX server gets the control of virtual machine's storage using VMFS. The virtual machine retains the same IP address after the migration. VMotion is the key enabling technology that allows VMware DRS to create a self-managing, highly optimized, and efficient virtual environment with built-in load balancing.

vSphere supports Enhanced VMotion Compatibility (EVC) using CPU features such as Intel FlexMigration and AMD-V Extended Migration technologies to allow migrations from older servers to newer servers.

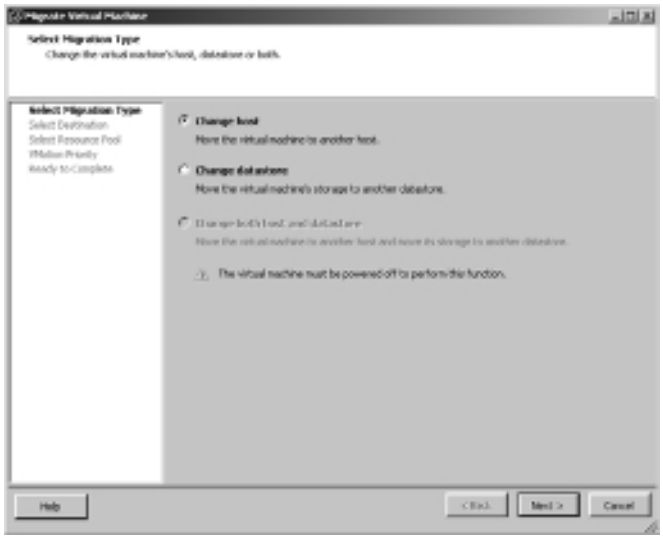

FIGURE 1.11
VMware vSphere VMotion types

VMware Storage VMotion

Storage VMotion allows you to migrate virtual machine disks for running virtual machines from one storage array to another. This avoids the need to schedule any service disruption or downtime because of planned storage maintenance or during storage migrations and upgrades. Storage VMotion was introduced in VMware Infrastructure 3.5, but it supported only the CLI and a limited number of storage protocols. Storage VMotion in vSphere is now integrated into vCenter and features several new capabilities:

- Includes full support for migrations from one storage vendor to another
- Works across NFS, Fibre Channel, and iSCSI storage protocols
- Includes an option to change the virtual disk format during a Storage VMotion session such as changing a thick to thin virtual disk format, or vice versa.

Storage VMotion in vSphere also features a more efficient migration process. It leverages a new and more efficient block copy mechanism in the VMkernel instead of using the virtual disk snapshot method found in the earlier releases of VMware products.

You can use Storage VMotion with thin provisioning to not only migrate a VM from one data store to another but also reclaim over-provisioned storage space during this process. Virtual machines often have guest OS virtual disks that are over-provisioned compared to what they truly need for their current storage requirements. With Storage VMotion, you can change the virtual disk format from thick to thin and reclaim this over-allocated but unused storage space (Figure 1.12).

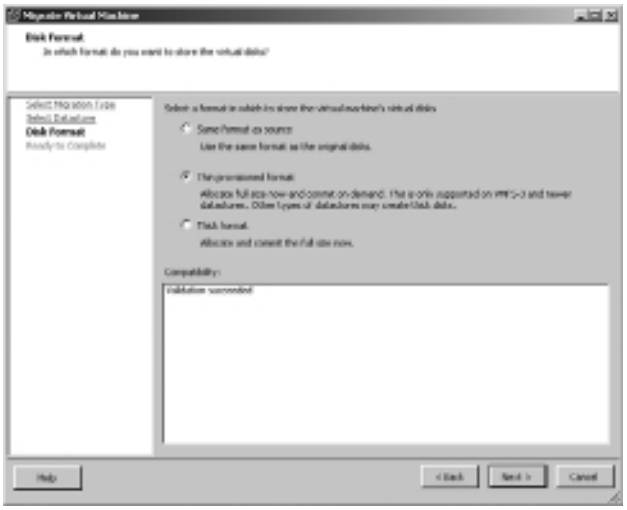

FIGURE 1.12
Storage VMotion, virtual disk format conversion

VMware High Availability

VMware High Availability (HA) extends high availability for any application running in a virtual machine, regardless of its operating system or underlying hardware configuration (Figure 1.13). VMware HA automatically detects physical server failures and operating system failures within virtual machines. If a failure is detected, VMware HA automatically restarts the affected virtual machine on a different ESX server in the resource pool. When restarting the affected virtual machines, you can also leverage DRS to take into account ESX host workload. The time it takes to reboot the virtual machine and restart the application is the only down-time that you will experience. VMware HA supports up to 32 ESX servers in a cluster. You can reserve a specific cluster capacity for failover or mark specific ESX servers as failover hosts. VMware HA can be a much simpler and cost-effective high-availability solution compared to traditional clustering solutions.

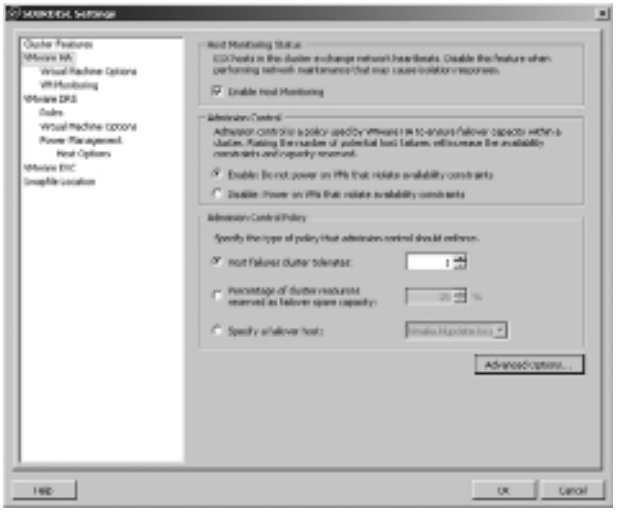

FIGURE 1.13
VMware HA settings

VMware Fault Tolerance

VMware Fault Tolerance (FT) provides zero downtime and zero data loss availability for all virtual machines during server hardware failures. When FT is enabled for a specific virtual machine, a secondary copy of that virtual machine runs in lockstep synchronization on another ESX host. This allows instantaneous, stateful failover between the two virtual machines and eliminates any disruption because of hardware failures on either host. FT does not support certain vSphere features and will display the warning shown in Figure 1.14 when you attempt to enable it.

FIGURE 1.14
VMware FT warning

VMware Data Recovery

VMware Data Recovery (VDR) provides a simple and easy-to-deploy disk-based backup and restore solution for protecting your virtual environment. VDR is deployed as a virtual appliance and does not require you to install any agents inside virtual machines. You back up and restore entire virtual machines and do not worry about the guest operating system or application running inside. It is integrated with VMware vCenter Server to provide a centralized user interface to manage your virtual machine backup and recovery jobs. VDR leverages virtual machine snapshots to make a precise copy in a short amount of time and also uses data *deduplication* technology to save on disk storage. VDR also supports Volume Shadow Copy Service (VSS) to enable consistent backups of virtual machines running Microsoft operating systems and applications. Depending on your business service-level agreements (SLAs) for Recovery Time Objective (RTO) and Recovery Point Objective (RPO), VDR can assist you in creating the perfect backup and recovery solution for your business organization. We discuss how you can use VDR in more detail in Chapter 2.

High Availability Clustering with Windows Server 2003 and 2008

vSphere supports Microsoft Cluster Service (MSCS) with Windows 2003 and Failover Clustering for Windows 2008. Clustering is supported with both 32-bit and 64-bit guests. Booting from SAN virtual machines is supported as well. vSphere also supports Majority Node Set clusters with application-level replication, such as Exchange 2007 Cluster Continuous Replication (CCR). The serial attached SCSI (SAS) virtual device introduced as part of vSphere virtual hardware version 7 is required to support Failover Clustering configuration for Windows Server 2008 and R2 versions. Later chapters in this book will cover setting up Failover Clustering for Windows Server 2008 R2 and will demonstrate how to use this device.

TIP VMware provides a separate guide to detail the setup for Failover Clustering and Microsoft Cluster Service. The guide is available at `www.vmware.com/pdf/vsphere4/r40_u1/vsp_40_u1_mscs.pdf`.

Security

Security services help you secure your virtual infrastructure from vulnerabilities. They also enable applications to enforce an appropriate level of security policies in an operationally efficient way. VMware vSphere includes the following security services:

VMware vShield Zones VMware vShield Zones is an application-aware firewall that can be integrated with VMware vCenter Server to enforce corporate security policies and ensure regulatory compliance at the application level in vSphere environments. It continuously monitors and controls the network traffic flowing to and from virtual machines in its inventory, while still maintaining trust and network segmentation of users and sensitive data.

VMware VMsafe VMware VMsafe provides an application programming interface (API) that security vendors can leverage to develop VMware-aware security products. VMsafe enables partners to build virtual appliance–based security solutions that can monitor and protect virtual machine CPU state, memory pages, network traffic, and disk files and any processes executing inside them.

In the next sections, we will provide more details about each of these security services.

VMware vShield Zones

VMware vShield Zones helps you protect the privacy and confidentiality of virtual machines and their data. vShield Zones builds an inventory of the operating systems, applications, and open ports within your virtual data center. You can then use this information to monitor and enforce network access to sensitive areas of the virtual data center, including the DMZ, or to servers with sensitive data that is subject to regulations such as PCI, SEC 1742a, or SOX compliance. It also allows you to build logical trust or organizational boundaries within existing vCenter Server deployments, while still leveraging the flexibility and availability of shared resource pools. You can then define security policies to bridge, firewall, or isolate traffic across these boundaries.

The components that make up vShield Zones environment are as follows:

vShield This is a virtual appliance located on each vSphere host and is used to inspect traffic flow and provide firewall protection.

vShield Manager This manages all of the distributed vShield instances by providing monitoring, centralized configuration, and software updates.

Once deployed, vShield sits in between your protected virtual machines and the external network interfaces. This allows vShield to intercept each network packet and enforce the policies that have been created for the particular security zone.

VMware VMsafe

VMware VMsafe enables an open security architecture with a set of APIs from VMware that gives security vendors the insight into the inherent properties of virtualization, similar to a hypervisor. Using this set of APIs, the security partners can develop virtual appliance–based security solutions that can monitor and protect virtual machine CPU state, memory pages, network traffic, and disk files and the processes executing inside them. Because these products will work in

conjunction with the virtualization layer, they are expected to provide higher levels of security to virtual machines compared to even physical servers. The following are examples of such VMsafe-based VMware-aware security products:

- A single antivirus virtual appliance that protects a vSphere host and all virtual machines running on it.
- A network security virtual appliance that protects network access for all virtual machines on a vSphere host.
- Security solutions that are aware of advanced VMware features such as DRS and vDS and that continue to protect virtual machines as they migrate using VMotion, Storage VMotion, or vDS network VMotion.

VMkernel Protection

vSphere introduces several security mechanisms to assure the integrity of the VMkernel and loaded modules as they reside on disk and in memory. It leverages disk integrity mechanisms to protect against malware, which might attempt to overwrite or modify VMkernel as it persists on disk. vSphere makes use of Trusted Platform Module (TPM), a hardware device embedded in servers, to protect the boot-up of the hypervisor.

VMkernel modules are digitally signed and validated during load time to ensure the authenticity and integrity of dynamically loaded code. VMkernel also uses memory integrity techniques at load time coupled with microprocessor capabilities to protect itself from common buffer-overflow attacks that are used to exploit running code. All these techniques are part of ongoing efforts to protect the hypervisor from common attacks and exploits and create a stronger barrier of protection around the hypervisor.

Scalability

Scalability services allow the vertical and horizontal scaling of virtual infrastructure while ensuring that the right amount of resources are allocated to applications without any disruption.

VMware ESX and ESXi scalability vSphere continues to extend the maximum supported configurations for the underlying physical hardware used by ESX and ESXi as well as virtual machines running on them. vSphere also makes the switch to 64-bit versions of VMkernel and Service Console for better scalability.

VMware DRS VMware DRS improves scalability in vSphere deployments in two ways. First, it automatically and continuously balances the workload across ESX servers within a cluster, making sure that no single virtual machine is bottlenecked on resources. Second, it provides a proportional resource allocation mechanism using shares, reservations, and limits so that you can dynamically control the resources used by a virtual machine without the need for any reboot.

Virtual machine hot-add support Virtual hardware version 7 in vSphere 4 introduces hot-add support for various virtual devices. The ability to add more resources to a virtual machine without powering it off can help you improve virtual machine scalability as needed. This functionality is supported only if the underlying guest operating system supports it.

VMware ESX and ESXi Scalability

vSphere 4 extends the scalability of the virtualization platform in several aspects. vSphere 4 supports servers with up to 64 logical CPUs and 1 TB of RAM. Consequently, vSphere can support up to a total of 512 virtual CPUs per single host. By increasing the number of virtual machines that can run on single host, vSphere can achieve a higher consolidation ratio compared to earlier versions of ESX.

At an individual virtual machine level, vSphere 4 adds support for up to eight virtual CPUs and 255 GB of memory. With these higher per virtual machine scalability limits, you can now virtualize more enterprise applications in your data center than before.

vSphere has also switched to 64-bit versions of VMkernel and Service Console. The 64-bit versions offer scalability benefits over 32-bit versions similar to other software applications. These benefits include the ability to address a lot more memory and better support for the newer 64-bit hardware.

VMware DRS

We discussed VMware DRS in detail in the "Infrastructure Services" section. Here we'll focus only on those DRS features that help improve scalability in vSphere deployments. DRS enables the ongoing dynamic load balancing of server resources within a cluster. This ensures that the applications are getting required resources all the time and no single virtual machine is bottlenecked on resources. The optimal virtual machine placement resulting from DRS load balancing can improve application scalability in a way that you cannot beat on a continuous basis.

DRS also allows you to allocate resources to virtual machines in proportion to the priorities established in the form of shares, reservations, and limits. For example, DRS will allocate proportionally higher shares of resources to virtual machines with higher shares or will guarantee fixed quantities of memory or CPU for virtual machines based on their reservations. You can dynamically shrink and grow virtual machine resource usage as needed without rebooting them.

Together, the dynamic load balancing and proportional resource allocation mechanisms can improve application scalability in vSphere deployments.

Virtual Machine Hot-Add Support

Virtual hardware version 7 introduced in vSphere 4 has hot-add support for various virtual devices:

- You can hot add CPU and memory to a virtual machine when needed without rebooting it (Figure 1.15).
- You can add or remove virtual storage and network devices from virtual machines without disruption.
- You can hot extend virtual disks for running virtual machines without any downtime.

This ability to dynamically add more resources to a virtual machine without powering it off can help you scale virtual machines as needed. However, this functionality needs to be supported by the guest operating system running inside the virtual machine. As a result, virtual machine hot-add support is currently supported for a limited number of operating systems only.

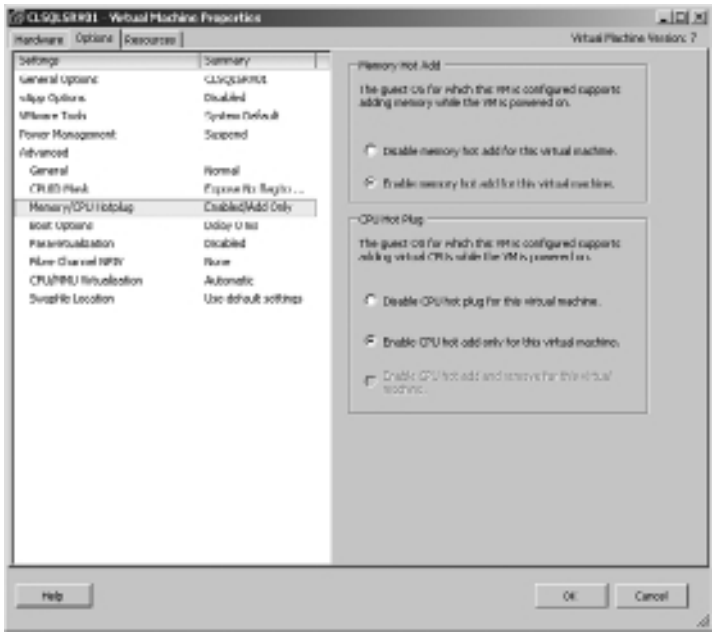

FIGURE 1.15
Virtual machine hot-add support for CPU and memory

Management Services

VMware vSphere management services include the tools and interfaces available to virtual infrastructure administrators to efficiently manage VMware vSphere environments. These management services include the following:

vCenter Server vCenter Server provides the central point of control for the administration and management of VMware vSphere environments. It acts as a single console to manage key IT processes for a virtualized data center such as monitoring, provisioning, patching, migration, capacity management, disaster recovery, and many other critical functions. VMware vCenter also integrates with existing industry-leading systems management tools from vendors such as BMC, CA, HP, and IBM/Tivoli for end-to-end physical and virtual management for a data center.

vCenter Orchestrator vCenter Orchestrator is a powerful orchestration engine integrated with vCenter. You can either use an out-of-the-box library of workflows or develop your own customized workflows to automate operational tasks.

vCenter Guided Consolidation vCenter Guided Consolidation provides a wizard-based process for first-time virtualization users to accelerate the server consolidation process. You can utilize this in smaller environments to discover physical servers, analyze their resource utilization, convert these physical servers to virtual machines, and place them on appropriate vSphere hosts.

vCenter Update Manager vCenter Update Manager automates the process of scanning and patching online VMware ESX hosts and selective Microsoft and Linux virtual machines to enforce compliance to patch standards. You can upgrade virtual machine hardware, VMware Tools, and virtual appliances as well as patch and update third-party software running on the virtual machines and virtual appliances.

vSphere Command-Line Interfaces vSphere provides two key command-line interfaces (CLIs) to automate common vSphere administration and management tasks: vSphere CLI and vSphere PowerCLI.

In the next sections, we will provide more details about each of these management services.

vCenter Server

VMware vCenter Server provides a single-pane-of-glass view for managing and administering the infrastructure and application services described in the previous sections. It provides visibility into every aspect of VMware vSphere environments and enables unified management of all the ESX/ESXi hosts and virtual machines in your data center from a single console along with aggregate performance monitoring of clusters, hosts, and virtual machines. VMware vCenter Server gives administrators insight into the status and configuration of clusters, hosts, virtual machines, storage, the guest OS, and other critical components of a virtual infrastructure—all from one place.

VMware vCenter Server lets you rapidly provision virtual machines and hosts using standardized templates. It creates a comprehensive map of the inventory and topology of your data center. It also controls access to virtual assets and functionality through fine-grained access controls, custom roles, permissions, and integration with existing Microsoft Active Directory authentication mechanisms. VMware vCenter Server also gives administrators control over key capabilities such as VMware VMotion, Distributed Resource Scheduler, High Availability, and Fault Tolerance.

In addition, VMware vCenter Server uses an open plug-in architecture to create a scalable and extensible platform. This allows VMware partners to integrate with vCenter Server to develop advanced management capabilities in areas such as capacity management, compliance management, business continuity, and storage monitoring. The vCenter Server APIs also allow customers to integrate physical and virtual management tools, using their choice of enterprise management tools to connect to vCenter Server.

vCenter Server 4 has been enhanced to include several new features such as the following:

- Simplified navigation and inventory search
- Host profiles
- Linked Mode
- Hardware monitoring with CIM
- Centralized licensing

We will provide more details about these features in the next sections.

Simplified Navigation and Inventory Search

VMware vCenter Server 4 has redesigned the landing page to provide users with an easy, one-click access to the vCenter inventory, vCenter solution plug-ins, and key administration and management tasks (Figure 1.16). In addition to the home page, vCenter Server also introduces a globally accessible navigation bar that makes it easy to switch between different parts of vSphere Client. The client application also remembers the view that was displayed when you log out of vSphere Client and will return you to that view when you next log in.

FIGURE 1.16
VMware vCenter inventory search

vCenter Server 4 also includes a new inventory search field to simplify locating virtual machines, hosts, data stores, networks, and folders based on the criteria you specify. You can perform simple searches based on keywords entered in the search field at the top right of vSphere Client. You can also perform advanced searches by specifying multiple search criteria—for example, virtual machines that need an update to VM Tools or data stores that have less than 10 GB of free space remaining.

HOST PROFILES

vCenter Server 4 introduces host profiles that can simplify host configuration management through user-defined configuration policies (Figure 1.17). You can use profile policies to eliminate any manual, single-host, ad hoc ESX host configuration and efficiently maintain configuration consistency and correctness across the entire data center.

Using host profile policies, you capture the blueprint of a known, validated "golden" configuration and use this to quickly configure networking, storage settings, security settings, and so on, across a large population of hosts. For example, you can quickly update DNS and NTP settings for several ESX hosts or configure multiple ESX hosts to use a new vNetwork Distributed Switch. Note that ESX hosts need to be in maintenance mode to apply a host profile.

You can also use host profile policies to monitor and report on compliance to standard host configuration settings across the data center (Figure 1.18). The profile compliance information displayed on the Profile Compliance tab depends upon the object selected in the vCenter inventory panel. For example, if you select a cluster in the vCenter inventory panel, you will see profile compliance information for all hosts within the selected cluster. vCenter generates a base set

of host compliance checks based on host configuration settings, such as networking, DNS, and NTP settings. In addition, vCenter performs certain built-in cluster compliance checks (shown in Table 1.1) even if a host profile is not attached to a cluster.

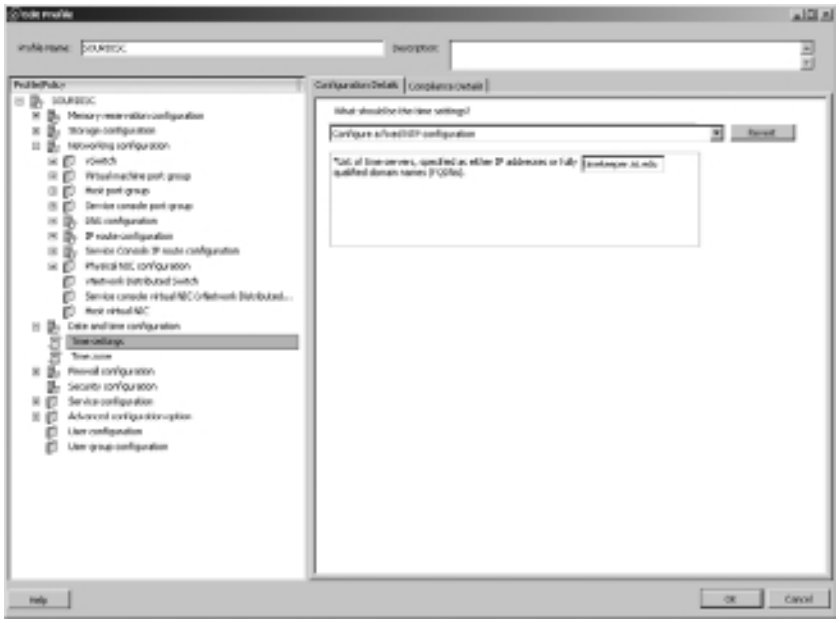

FIGURE 1.17
Host profiles' NTP configuration

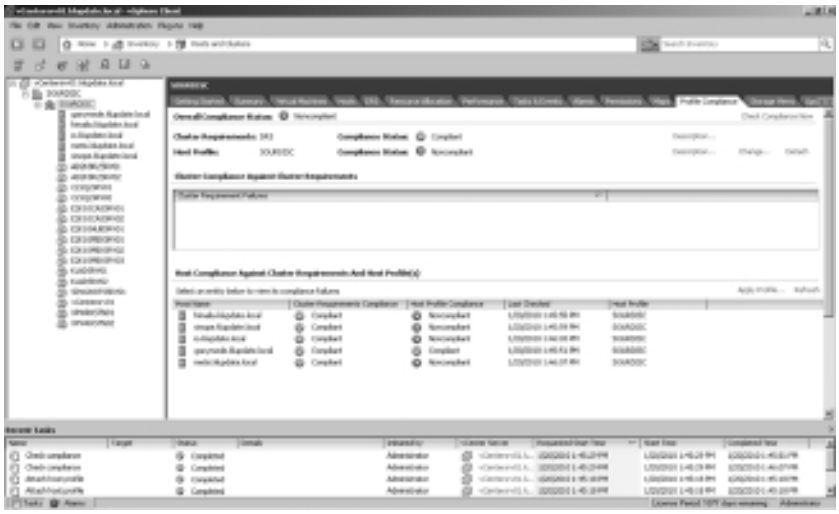

FIGURE 1.18
Profile compliance

TABLE 1.1: Built-in Cluster Compliance Checks

CLUSTER REQUIREMENT	CLUSTER COMPLIANCE CHECK
VMware DRS	Validate that VMotion NIC speed is at least 1000 Mbps.
	Validate that VMotion is enabled.
	Validate that at least one shared data store exists.
VMware DPM	Validate that power management is supported on the host.
VMware HA/VMware FT	Validate that FT logging is enabled.
	Validate that FT logging NIC speed is at least 1000 Mbps.
	Validate that all the hosts in the cluster have the same build for FT.
	Validate that the host hardware supports FT.

vCenter Server Linked Mode

A single instance of vCenter Server 4 can manage up to 300 hosts and 3,000 virtual machines. This limit is not enough for organizations ramping up their virtualization projects to span an entire data center. To meet these higher scalability requirements, vCenter Server 4 introduces a feature called Linked Mode, where multiple vCenter Server systems can be linked together and monitored from a single vSphere Client session. Using the Linked Mode feature, you can manage up to 1,000 hosts and 10,000 virtual machines across 10 vCenter Server instances.

Linked Mode leverages Microsoft Active Directory Application Mode (ADAM), an implementation of Lightweight Directory Access Protocol (LDAP), to store and synchronize data across multiple vCenter Server instances. ADAM is installed automatically as part of the vCenter Server 4 installation. The ADAM instances in a group use peer-to-peer networking to replicate the following information for each vCenter instance to the LDAP directory:

- Connection information (IP and ports)
- Certificates and thumbprints
- Licensing information
- User roles and permissions

When vCenter Servers are connected in Linked Mode, you can do the following:

- Log in simultaneously to all vCenter Servers using a valid credential
- Search the inventories of all the vCenter Servers in the group
- View the inventories of all the vCenter Servers in the group in a single inventory view

When you log in vCenter using Linked Mode, you will see the inventory of all vCenter instances at once. The inventory tree on the left side will show each vCenter instance at the top level. You can then use the +/- indicator to expand or collapse the inventory tree to, respectively, show or hide lower-level objects such as the data stores, folders, clusters, hosts, and so on, for

any vCenter instance. Note that a user needs to have valid permissions to be able to see a vCenter Server instance.

The vCenter Linked Mode feature's visibility across multiple vCenter instances applies only to view and search operations. Other operations are confined within a single vCenter inventory. For example, you cannot drag and drop a host between vCenter instances or a virtual machine between hosts on two different vCenter instances.

You can join a vCenter instance to a Linked Mode group at the time of installation or afterward by modifying an existing deployment. Both of these methods are described in the vSphere Installation Guide.

Hardware Monitoring with CIM

vCenter Sever 4 provides a Hardware Status plug-in to monitor and manage the health of your VMware ESX servers, including key components such as fans, system boards, and power supplies (Figure 1.19). The vCenter Hardware Status plug-in uses the industry-standard Common Information Model (CIM) interface to display this health information. It implements the System Management Architecture for Server Hardware (SMASH) profiles defined by the Distributed Management Task Force (DMTF).

Figure 1.19
vSphere hosts' hardware status

The vCenter Hardware Status plug-in provides you with an integrated, centralized view of both the physical and virtual assets in your vSphere environment. The plug-in is also integrated into the vCenter alarm interface so you can be alerted when hardware failures occur. You can also trigger automated alarm workflows based on these vCenter alarms to preempt and remedy hardware problems. For example, if vCenter detects that the host temperature is getting too hot, it could trigger an alarm action that automatically puts the VMware ESX host into maintenance mode (migrating the VMs off the host using VMware DRS in the process) to allow the server to cool down.

Centralized Licensing

In vSphere 4, VMware has redesigned the licensing facility. License reporting and management are centralized into vCenter Server 4. Unlike VMware Infrastructure 3, there is no separate license server that must be installed and monitored. Using the vCenter Server interface, you can centrally assign VMware vSphere licenses, report on license usage, and monitor for license state and compliance. If you upgrade all of your hosts, you no longer need a license server or host license files. All product and feature licenses use 25-character license keys that you can manage and monitor from vCenter Server 4.

When vCenter Server 4 assigns a license key to a vSphere host, the license key is copied to the host and saved in a persistent format. In the event that the host becomes disconnected from vCenter Server 4, the license key remains active on the host indefinitely, even after a host reboot. Only a deliberate licensing operation by the user can remove or replace a host license key. VMware recommends assigning all VMware vSphere licenses centrally from vCenter Server 4; however, you can also assign license keys directly to individual hosts.

For more information on VMware vSphere licensing, go to `www.vmware.com/support/licensing.html`.

vCenter Orchestrator

vCenter Server 4 includes vCenter Orchestrator, a workflow orchestration engine to help you automate management tasks for vSphere either by using out-of-the-box workflows or by assembling workflows using an easy drag-and-drop interface.

Workflows are reusable building blocks that combine actions, decisions, and results that, when performed in a particular order, complete a specific task or process in a virtual environment. You can leverage the out-of-the-box library of extensible workflows provided with vCenter Orchestrator to create and execute automated, configurable processes to manage your infrastructure. Or you can use vCenter Orchestrator to capture best practices within your data center for tasks, such as provisioning virtual machines, backing up, and performing regular maintenance, and turn them into workflows. vCenter Orchestrator also provides access to all operations in the vCenter Server API. You can integrate all these operations into your automated processes or workflows. This will help you achieve consistency, standardization, and overall compliance with existing IT policies in your virtual environment.

Orchestrator also allows you to integrate with other third-party management and administration solutions through its open plug-in architecture.

vCenter Guided Consolidation

vCenter Guided Consolidation provides a wizard-based process for first-time virtualization users to accelerate the server consolidation process. The Guided Consolidation service is provided as a modular plug-in to vCenter Server. To reduce the overhead on the vCenter Server, the Guided Consolidation service can also be installed on a different system than vCenter Server 4.

As the name suggests, Guided Consolidation guides you through all the steps of the consolidation process:

- Automatic discovery of physical servers
- Performance analysis to identify virtualization candidates

- Physical-to-virtual conversion
- Intelligent placement on the right vSphere host

You can find and analyze physical systems using Guided Consolidation in several ways:

- You can automatically scan active domains daily.
- You can manually add systems by entering a computer name, a singular or range of IP addresses, or a filename.

Guided Consolidation is recommended for smaller environments, up to 100 concurrent physical machines at any given time. For larger environments, you should use VMware Capacity Planner.

vCenter Update Manager

vCenter Update Manager automates the process of patch management in vSphere environments and helps you enforce compliance to patch standards. The vCenter Update Manager can help you do the following:

- Patch and upgrade VMware ESX/ESXi hosts
- Apply patches to Windows and certain versions of Linux guest operating systems in virtual machines
- Upgrade VMware Tools and virtual hardware for virtual machines
- Patch and upgrade virtual appliances

VMware Update Manager integrates with the Shavlik patch management technology to provide these patching capabilities. It also allows you to download patches from a remote server and stage them to a local server. Staging patches on local server can be very helpful when patching several hosts and virtual machines.

Update Manager also allows you to create update baselines and baseline groups that represent a set of updates. You can then review the Update Manager dashboard to compare hosts and virtual machines against these baselines. If the host or virtual machine is not compliant, then it can be easily remediated. You can use this to ensure consistency across ESX hosts and virtual machines in your vSphere environment.

In Chapter 2, we discuss in more detail how you can use Update Manager to patch Windows Server 2008 virtual machines and comply with your patch management policies.

vSphere Management Assistant

With vSphere, VMware also ships vSphere Management Assistant (vMA), a virtual machine that includes the vSphere CLI, vSphere SDK for Perl, and other prepackaged software. You can use this prebuilt virtual machine to run agents and scripts to manage ESX/ ESXi and vCenter Server systems. Some folks refer to vMA as the missing service console for ESXi. A key functionality of vMA includes noninteractive login, which allows you to use the vSphere CLI without having to explicitly authenticate each time. vMA can also collect ESX/ESXi and vCenter Server logs and store the information for analysis. vMA can also host third-party agents for added management functionality.

vSphere Host Update Utility

vSphere also includes the vSphere Host Update Utility to enable you to remotely upgrade ESX hosts (versions 3.0 and newer) to ESX 4.0. It upgrades the virtual machine kernel (VMkernel) and the Service Console, if present, but does not upgrade VMFS data stores or virtual machine guest operating systems. vSphere Host Update Utility comes as a stand-alone Microsoft Windows application that provides real-time status of a remote upgrade as well as allows you to specify custom post-upgrade scripts. The utility performs an automated host compatibility check as a pre-upgrade step to verify that each host is compatible with ESX 4.0/ESXi 4.0. The Host Update utility also includes a new rollback feature to back out failed upgrades. It is recommended for smaller deployments with fewer than 10 ESX/ESXi hosts, without vCenter Server or Update Manager.

vSphere Command-Line Interfaces: vCLI and PowerCLI

vSphere provides two key command-line interfaces to automate vSphere administration and management tasks:

vSphere CLI The VMware vSphere CLI (vCLI) provides command-line interface tools for managing vSphere ESX and ESXi servers. You can use this with ESX/ESXi and vCenter Server 4 systems as well as with ESX/ESXi versions 3.5 Update 2 and newer. vCLI includes more than 30 command-line utilities to help you provision, configure, and maintain your ESX and ESXi hosts.

vSphere PowerCLI Based on Microsoft PowerShell technology, VMware vSphere PowerCLI is a powerful command-line tool that lets you automate all aspects of vSphere management, including network, storage, ESX, VM, guest OS, and more. PowerCLI is distributed as a Windows PowerShell snap-in, and it includes more than 150 PowerShell cmdlets, along with documentation and samples.

VMware vApp

VMware vSphere includes support for vApp, a single logical entity that encapsulates an *n*-tier application comprising one or more virtual machines. vApp uses the industry-standard Open Virtualization Format (OVF) to specify and encapsulate all components of a multitier application as well as the operational policies and service levels associated with it. Along with virtual machines, vApps also captures the dependencies between these virtual machines as well as the resource allocations associated with these virtual machines. Defining all this information in one logical entity makes it very convenient, often in a single step, to power off/on, clone, deploy, and monitor an entire application. vCenter Server can create and run vApps, as well as import and export them in compliance with the OVF 1.0 standard.

VMware Management Solutions

Complementing vCenter Server and its management services, VMware also provides several virtualization management solutions. Here are a few of these products:

- VMware vCenter AppSpeed to monitor and troubleshoot the performance of applications virtualized in a VMware vSphere environment
- VMware vCenter CapacityIQ to perform capacity sizing and management of virtual machines, clusters, and data centers

- VMware vCenter Chargeback to automate tracking and chargeback for the cost of IT services
- VMware vCenter Lab Manager to create an on-demand lab infrastructure
- VMware vCenter Lifecycle Manager to automate and manage virtual machine provisioning in your virtual infrastructure
- VMware vCenter Site Recovery Manager to automate the setup, testing, and failover of disaster recovery processes

In this section, we provide a brief overview for each of these solutions.

VMware vCenter AppSpeed

VMware vCenter AppSpeed enables you to monitor application-level latency, throughput, and transaction rate and correlate the data with performance indicators at the infrastructure level. It can monitor most multitier applications, including database traffic (Microsoft SQL, MySQL, Oracle), any HTTP or HTTPS traffic (most application servers and web servers), and Microsoft Exchange Server traffic. It does this without installing an agent or requiring login credentials to the applications. AppSpeed is shipped as a virtual appliance, so it can be configured and deployed with minimal effort. When deployed, AppSpeed monitors application traffic through virtual switches and automatically builds application topology and maps interdependencies between application components. AppSpeed can detect the slowest or most used application transactions. It also allows you to drill down into individual transactions, as well as point out transaction dependencies across components in a multitier application stack. This can help you quickly isolate and remediate application performance problems in vSphere deployments. With AppSpeed, you can also set up ongoing performance service-level monitoring based on thresholds for transaction latency. You can either set this threshold manually or let AppSpeed automatically select the values based on the historical information it has.

We discuss how you can deploy and use AppSpeed in Chapter 5. For more details, you can also check the vCenter AppSpeed product page at `www.vmware.com/products/vcenter-appspeed/`.

VMware vCenter CapacityIQ

VMware vCenter CapacityIQ adds capacity management capabilities to your vSphere deployment. It enables ongoing, real-time tracking of capacity utilization in your environment and analyzes this information to perform capacity trending and forecasting. With CapacityIQ, you can tell how long your capacity will last and when it's time to add more capacity. CapacityIQ can also help you optimize current capacity usage and size your infrastructure correctly by identifying idle, powered-off, over-provisioned, and under-provisioned virtual machines. CapacityIQ also provides "what-if" modeling to help you understand the impact of adding or removing virtual machines or vSphere hosts to your deployment. You can model the virtual machines or vSphere hosts based on expected usage or even base them on existing VMs. CapacityIQ is shipped as a virtual appliance, so it can be configured and deployed with minimal effort.

For more details, check the vCenter CapacityIQ product page at `www.vmware.com/products/vcenter-capacityiq/`.

VMware vCenter Chargeback

VMware vCenter Chargeback enables you to track how virtual infrastructure is being used by different departments in your organization, correlate it to their cost centers, and finally charge them for the consumption. Chargeback lets you create variety of cost models—fixed, reservation, usage, or a hybrid cost model—that measure cost on a per-VM basis. These cost models can be applied at the business level in a chargeback hierarchy, so there is no need to create cost models for each VM. It supports allocation-based or utilization-based chargeback, or a hybrid of both. You can leverage this to set up tiered cost models that vary across cost centers, or types of use requirements. For example, a production environment may cost more than a development environment. Similarly, you will need to pay a premium for a highly available virtual machine with HA or FT enabled. Using Chargeback's detailed reports, you can either bill or simply "show" to the business groups the IT cost of their virtual infrastructure usage.

For more details, check the vCenter Chargeback product page at `www.vmware.com/products/vcenter-chargeback/`.

VMware vCenter Lab Manager

VMware vCenter Lab Manager allows to you deploy multiple virtual machines for a development or test environment from a library of existing configurations using a self-service portal. It has an intuitive user interface that is simple enough for non-IT users but also includes advanced features such system quotas and access control to manage resource allocation. You can control multiple machines as a single atomic unit and deploy them to a development or test or staging environment. It is also widely used in training and education to provision lab environments, in support and help-desk groups for troubleshooting, and in sales and marketing for live product demos and for product evaluations.

For more details, check the vCenter Lab Manager product page at `www.vmware.com/products/labmanager/`.

VMware vCenter Lifecycle Manager

VMware vCenter Lifecycle Manager provides you with better control over virtual machine provisioning and deployment in your vSphere deployment. It provides workflows built on top of vCenter Orchestrator to automate the entire virtual machine life-cycle management from provisioning to decommissioning. Lifecycle Manager provides a centralized web portal to enable self-service provisioning of virtual machines. The end users choose from a catalog of virtual machine configurations and are allowed to request or create virtual machines based on user policies and access control rights. Lifecycle Manager then automatically tracks the virtual machine throughout its life.

For more details, check the vCenter Lifecycle Manager product page at `www.vmware.com/products/lifecycle-manager/`.

VMware vCenter Site Recovery Manager

VMware vCenter Site Recovery Manager addresses the disaster recovery challenges for the entire data center. By leveraging virtualization and replication capabilities of storage, it enables more affordable disaster recovery protection than previously possible. With vCenter Site Recovery Manager, you can build multiple disaster recovery workflows within your vCenter

implementation to automate your failover plans. You can also test these workflows completely so that you know the recovery will work if and when you need it. It utilizes storage replication to ensure that data is successfully and accurately transferred to the failover site. Because of virtualization, you do not have to maintain strict hardware compatibility or map servers one to one with your primary site environment.

For more details, check the vCenter Site Recovery Manager product page at `www.vmware.com/products/site-recovery-manager/`.

VMware vSphere Editions

VMware provides vSphere in two separate offerings: one targeted for small businesses and the other for midsize and enterprise businesses. Each of these offerings includes different editions of vSphere that are tiered, based on the features included. In this section, we will provide a quick overview of these offerings.

vSphere for Small Businesses

For the small businesses, these are the vSphere editions available:

- VMware vSphere Essentials
- VMware vSphere Essentials Plus

Both these editions are limited to three physical servers and allow you to create virtual machines with up to four virtual CPUs. Unlike the vSphere editions for midsize and enterprise businesses, these editions also include the license for VMware vCenter Server for Essentials to help you manage the virtual environment. The Essentials Plus edition adds VMware HA and VMware Data Recovery to the Essential edition to provide the high-availability and data protection features. Note that these editions for small businesses are self-contained packages that cannot be combined with other VMware vSphere editions.

If you want to virtualize only one physical server, you can use VMware ESXi, which is the free version of the VMware ESX hypervisor with a smaller footprint. You will need to upgrade to other vSphere editions to leverage any additional vSphere features or use vCenter Server for management.

To address the needs of smaller deployments at retail and branch offices, VMware provides special editions of the vSphere Essentials packages:

- VMware vSphere Essentials for Retail and Branch Offices
- VMware vSphere Essentials Plus for Retail and Branch Offices

These retail and branch office editions allow you to upgrade the vCenter Server for Essentials edition to vCenter Server Standard edition. This Standard edition adds support for vCenter Orchestrator and vCenter Linked Mode features. The Essentials Plus edition adds VMware HA and VMware Data Recovery to the Essential edition to provide the high-availability and data protection features.

You can find the detailed comparison between these small business editions at `www.vmware.com/products/vsphere/buy/small_business_editions_comparison.html`.

vSphere for Midsize and Enterprise Businesses

VMware provides the following vSphere editions for mid-size and enterprise businesses:

- VMware vSphere Standard
- VMware vSphere Advanced
- VMware vSphere Enterprise
- VMware vSphere Enterprise Plus

These editions are tiered based on the vSphere features included. You can find the detailed comparison between these midsize and enterprise business editions at `www.vmware.com/products/vsphere/buy/editions_comparison.html`.

VMware vCenter Server is a required component for vSphere deployment and has to be purchased separately for these midsize and enterprise editions of vSphere. vCenter Server is also available in two editions:

- VMware vCenter Server Foundation
- VMware vCenter Server Standard

The Foundation edition is limited to managing three vSphere hosts. The Standard edition removes that limit and also adds support for vCenter Orchestrator and vCenter Linked Mode features.

TIP VMware provides a vSphere Purchase Advisor that can help you figure out the right edition of vSphere for your environment. The tool is available at `www.vmware.com/products/vsphere/purchase-advisor/`.

vSphere Compatibility Guides

VMware maintains strict compatibility lists for its virtualization platform. This is one of the key reasons why vSphere is rated as a mainframe-like, highly reliable, and resilient virtualization platform. vSphere is certified across the complete stack of servers, storage, operating systems, and software applications. VMware works closely with the various hardware providers to certify new hardware with vSphere. You can check the compatibility guides at `www.vmware.com/resources/compatibility`. Because of ongoing certifications, these compatibility guides are updated from time to time. You can leverage this information to plan and deploy hardware and operating systems across the entire data center. The online compatibility guides allow you to search for vSphere compatibility in the following categories:

- Server systems
- Storage/SAN
- I/O devices
- Guest/host operating systems
- VMware View-compatible devices

The systems, storage, and I/O devices compatibility guides list the hardware that's compatible with various versions of VMware ESX and ESXi. The guest and host operating system compatibility guides list the operating systems supported by various VMware products such as ESX/ESXi, Workstation, Fusion, and Server.

Summary

With the release of vSphere, VMware continues to expand the capabilities of its successful VMware Infrastructure virtualization platform. VMware classifies the vSphere platform functionality into three groups of services:

- Infrastructure services
- Application services
- Management services

Each of these services further comprises several components that together make vSphere a reliable and scalable virtualization platform. Although infrastructure and application services are part of vSphere, the management services are provided by VMware vCenter Server. In this chapter, we provided a high-level overview of these services and details on the different components that make these services. Understanding the capabilities of the underlying virtualization platform is the key to subsequent successful application deployment. You can leverage this information when we explain how to successfully virtualize the tier-1 Microsoft server applications in the next chapters.

Chapter 2

Virtualizing Microsoft Server Applications

According to an IDC press release dated February 2010 (`www.idc.com/getdoc.jsp?containerId=prUS22224510`), the Windows operating system and the server applications based on it are the market leaders in enterprise data centers. In 2009, x86 servers accounted for $23.7 billion in worldwide revenue and 6.4 million units of worldwide server shipments. This represented more than 55 percent of all server revenue and more than 96 percent of all server units shipped worldwide.

Windows-based servers account for about a 70 percent share of the x86 server revenue as well as server unit shipments. As customers ramp up their virtualization efforts, these Windows-based servers are being deployed in virtual environments at a rapid pace. Because VMware vSphere is the most preferred virtualization platform today, customers need guidance and best-practice recommendations for deploying their Windows server applications, such as Windows Server 2008, Active Directory, IIS, SQL Server, Exchange Server, SharePoint Server, and Remote Desktop Services on VMware vSphere.

VMware has published extensive guidance for virtualizing business-critical applications as well. This information is available on VMware's website at `www.vmware.com/solutions/business-critical-apps/`. The information extends beyond Microsoft server applications such as Exchange Server, SQL Server, and SharePoint Server, and also includes other enterprise server applications such as Oracle, SAP, and other industry verticals.

In this chapter, we discuss the topics that are common across these Windows server technologies and applications when virtualizing with VMware vSphere, such as virtualization technologies, hardware issues, virtual deployment and deployment strategies, licensing, data backups, patches, and updates.

The New Era of Virtualization

There's very little doubt that Microsoft has done a great job at infiltrating the data center. Walk into any enterprise (well, maybe not Oracle or Red Hat), and you will no doubt find hints of Exchange Server, SQL Server, and SharePoint Server among others. These are considered Microsoft's tier-one applications. In other words, they are applications that businesses consider a critical part of running their day-to-day operations. Ensuring these applications perform optimally, are highly available, and are deployed in a manner that is within the bounds of support statements are the key concerns when deploying any of these applications.

Over the past few years, there's been a change to the landscape upon which these applications are deployed. Virtualization is no longer seen as only a development or testing tool; companies

are increasingly adopting the practice of virtualizing their data centers, and vSphere is at the center of this revolution. This upward trend of virtualizing the data center had mainly focused on the low-hanging fruit, or those less-than-critical applications that were thought of as easy to virtualize. With the low-hanging fruit out of the way, tier-one applications can now get the attention they deserve, and there couldn't be a better time.

There's no doubt that vSphere is the best-of-breed virtualization platform for deploying critical Microsoft applications, but a common concern for VMware customers has been the supportability of these solutions in a virtualized environment. These concerns can be put to rest because Microsoft now fully supports an array of virtualized deployments under the Server Virtualization Validation Program. The SVVP program is discussed in detail later in this chapter. In addition to the SVVP program helping push the virtualization of business-critical apps, VMware continues to provide performance studies and guidance around these applications; see `www.vmware.com/solutions/business-critical-apps/`.

These studies and white papers, along with the guidance provided in the chapters to follow, will help you in deploying the applications your business depends on as they were intended and in a supported fashion.

Virtualization Technologies

The virtualization technologies used to deploy your applications will greatly depend on the use case of your environment. For industry-standard x86 systems, the two most common virtualization technologies are hosted virtualization and bare-metal virtualization. In this section, we will discuss the differences between these two technologies.

Hosted Virtualization

Hosted virtualization was the first x86 virtualization technology widely available to the public to virtualize an x86 PC or server and its available hardware. This type of virtualization is known as a *type 2* or *hosted* hypervisor because it runs on top of the Windows or Linux host operating system. As shown in Figure 2.1, the hosted hypervisor allows you to run virtual machines as applications alongside other software on the host machine.

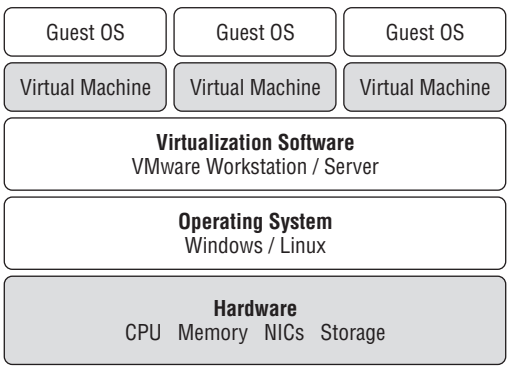

FIGURE 2.1
Hosted virtualization products run atop a host operating system.

VMware Workstation, Fusion, and Server are examples of hosted virtualization products.

The host operating system controls the access to the physical resources. The guest OS inside the virtual machine has to traverse through the host operating system layer before it can access these physical resources, thus introducing significant performance degradation within the virtualized guest OS environment. This type of virtualization is suitable for testing and developmental environments where developers can test their code in self-contained test and developmental environments.

This technology is not well suited for production use cases, and we strongly recommend that you do not deploy any application that must maintain data integrity or would interrupt your daily business operations because of possible data corruption in the event of virtual machine failure.

Bare-Metal Virtualization

The second type of hypervisor is known as *type 1 native* or *bare-metal*. This is commonly implemented as software running directly on top of the hardware with direct access and control of the hardware's resources. Since it has direct access to the hardware resources rather than going through an operating system (Figure 2.2), the hypervisor is more efficient than a hosted architecture and delivers greater scalability, robustness, and performance.

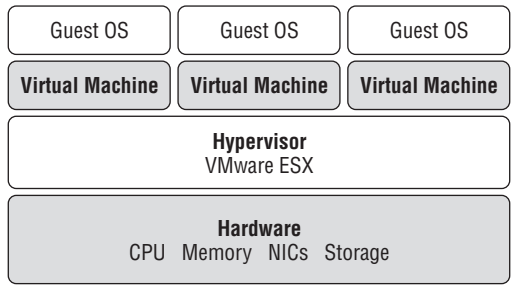

FIGURE 2.2 Bare-metal hypervisors have direct access to the physical hardware, providing better performance than hosted virtualization products.

VMware ESX is an example of the bare-metal hypervisor architecture. We strongly recommend that you use a bare-metal hypervisor such as VMware ESX (or its free edition, ESXi) when deploying server applications in a virtual environment.

Because the hypervisor directly controls the physical resources and the virtual machine hardware is abstracted and captured in a set of files, this type of virtualization also enables interesting use cases, such VMotion and VMware HA, which can significantly benefit your application availability.

WARNING The use of any VMware or Microsoft hosted virtualization technologies in production environment use case scenarios is not recommended or supported. Please see the VMware and Microsoft websites for more information on their hosted virtualization products.

Hardware Considerations

The x86 architecture was not designed to support virtualization. More specifically, these processors have 17 CPU instructions that require special privileges and can result in operation system instability when virtualized. VMware invented the "binary translation" technique to enable virtualization for x86 servers. To handle the problematic instructions, the binary translation technique "traps" these instructions as they are generated and converts them into safe instructions that can be virtualized, while allowing all other instructions to be executed without intervention. Binary translation provides broad compatibility, allowing you to use virtualization on practically any x86 server without the need to modify the guest operating system.

Virtualization technology has made significant advances in the past few years. Both AMD and Intel have embraced virtualization and developed new processor features to overcome the initial shortcomings in x86 architecture. First-generation enhancements included processors with Intel Virtualization Technology (VT-x) and AMD's AMD-V technologies that enabled CPU instruction set virtualization inside processors. For most workloads, the binary translation technique outperformed the first-generation hardware assist implementations.

Newer processors from AMD and Intel now also include hardware support for memory management unit (MMU) virtualization to alleviate this situation. Hardware-assisted memory management capabilities from AMD and Intel are called Rapid Virtualization Indexing (RVI) and Extended Page Tables (EPT), respectively. This support consists of a second level of page tables implemented in hardware to manage guest physical to machine memory address translations. With this addition, this second-generation hardware assist virtualization is usually faster compared to the binary translation technique.

When considering hardware for your virtualization projects, we recommend you use the newer-generation processors from AMD and Intel. Since the virtualization is offloaded to hardware, virtualized applications perform much better with less virtualization overhead. In fact, with vSphere, VMware prefers hardware assist virtualization over binary translation when creating virtual machines. VMware ESX 4.0 introduces support for the Intel Xeon processors that support EPT. Support for AMD RVI has existed since ESX 3.5.

Planning Virtual Deployment

Virtualizing Microsoft infrastructure and applications is not that much different from designing and implementing these same applications in a physical environment. Most if not all of the same principles you would apply in the physical design and implementation will apply in the virtual environment. The significant difference comes when you begin deploying multitiered applications such as SharePoint and Exchange. SharePoint, for example, requires database, application, and web layers. Traditionally this would require three separate physical servers, which given the power in today's enterprise-class servers would mean many wasted computing resources. Virtualization allows for taking these enterprise-class servers and creating independent virtual machines each with their own sets of resources. This is the very essence of virtualization: you can maximize your computing investment while minimizing your physical and environmental footprint. When planning for a virtual deployment, bigger is better. More processing and memory power means that more virtual machines can coexist, thus making efficient use of your computing resources.

You must take into account your business requirements regarding the following:

- Access security
- Performance
- Infrastructure and application maintenance
- Service-level agreements (SLAs)
- Disaster recovery and business continuity

In all instances, these are the same considerations you would make to ensure that your deployment projects are successful and supportable.

Common reasons that virtualization implementations fail are a lack of planning or understanding of the workload to be virtualized. This can lead to the deployment of insufficient hardware resources, an improper storage configuration for a specific application type, or a failure to plan for the aforementioned bullet points.

When planning a virtualized implementation of a Microsoft application or infrastructure product, you should always follow Microsoft best practices for the specific application or applications that you are designing and implementing. Microsoft has available resources for all of its applications, making it easy to design and implement those applications.

When planning for virtualization, ensure that you follow the VMware best practices for virtual infrastructure design and implementation. This includes the following:

- Deploying on the newest virtualization capable hardware
- Planning for availability
- Properly sizing the hosts for the applications that will be deployed on the host (including CPU count and memory)
- Properly designing the network (virtual as well as physical)
- Designing and using the proper storage

Understanding Operational Requirements

When planning a migration to a virtual environment, you should also understand other deployment and operational requirements, such as those for creating, staging, and provisioning virtual machines. How can you leverage virtualization to ease deployment? The deployment of virtual machines running Windows operating systems are streamlined in vSphere. The integration of a SysPrep-like interface allows for the quick deployment of a preconfigured "golden" image. This will be covered in detail in Chapter 3. Deployment is one small aspect of operational procedures; patching the operating system as well as backing up and restoring the application data is a common concern no matter the deployment scenario. Luckily, vSphere has integrated both of these functions into a single pane of glass, the vCenter console. Patches for the Windows operating systems are easily deployed using vCenter Update Manager, and backups of Windows VMs utilize the VSS framework when VMware Data Recovery is used. Both of these features are discussed in more detail later in the chapter.

You start planning the migration to a virtual environment by first understanding the workloads you plan on virtualizing. This can be done using Perfmon data, application-specific tools like the Exchange Server Profile Analyzer, or the Service Center Operations Manager, which contains specific counters for the Microsoft applications you are considering virtualizing. With an understanding of the performance requirements, you can proceed to rigorous piloting and testing by virtualizing a small set of server applications at a time. After the pilot phase has completed and you have had a chance to validate the test results, increase the pilot to include additional server applications and more challenging workloads. This phased approach allows you to ease into a virtual environment while realizing the benefits and capabilities of vSphere.

TIP For more information on Microsoft Solution Accelerators for Microsoft infrastructure and applications, please visit the Microsoft Solution Accelerators website at `http://technet.microsoft.com/solutionaccelerators`.

Assessing Your IT Infrastructure Using VMware Capacity Planner

When planning a migration to a virtual environment, you are more likely to succeed if you understand existing physical IT infrastructure and know the resource utilization. A virtualization assessment exercise that takes into account this baseline information and explores various server consolidation opportunities within your existing physical IT infrastructure can be helpful to set appropriate expectations. We recommend you do this due diligence properly so that you are better prepared to roll out your VMware vSphere deployment. VMware provides a tool, Capacity Planner that can help with this virtualization assessment exercise.

The Capacity Planner tool collects up to 30 days of system and performance data for existing Windows servers in your environment. The tool then analyzes the data to recommend which servers can be virtualized. Capacity Planner can also model various server consolidation scenarios based on inputs such as hardware configuration, processor models, and so on. The reports also include a preliminary total cost of ownership analysis of the sampled environment and recommended virtual deployment. You can access all this information from the web-based Capacity Planner Dashboard, as shown in Figure 2.3.

FIGURE 2.3
VMware Capacity Planner Dashboard

VMware Capacity Planner services are provided by both VMware as well as VMware partners. Visit the VMware Capacity Planner website for more information: `www.vmware.com/products/capacity-planner/`.

Licensing Considerations

Licensing can be a complex topic based on customer requirements, so it is generally recommended that you consult a licensing specialist from your reseller or a Microsoft Regional Licensing Specialist to ensure that you get the right answers to your licensing questions as well as to ensure that you operate your infrastructure and business-critical applications within Microsoft licensing compliance. Microsoft has updated its software licensing policies for server applications to make it more virtualization friendly.

Most Microsoft desktop, server, and applications use the server and client access license (CAL) model. Some products that follow the server/CAL model are as follows:

- Windows Server
- Exchange Server
- SharePoint Server
- ISA Server
- Groove Server

These are just a few of the Microsoft applications that use the server/CAL licensing model, which will benefit from the new licensing.

For customers planning to leverage VMotion, new Microsoft policies allow unlimited mobility of a large number of applications (including SQL Server 2008) under defined conditions. Microsoft has waived the limitation on short-term (90 days or less) license reassignment and allows customers to reassign licenses freely across servers within a server farm. With the changes, both licenses and software can move more freely across servers in a server farm, potentially reducing the number of licenses needed to support your workloads.

NOTE It is important to point out that these changes may not apply to software licenses for the Windows Server operating system or Remote Desktop Services CALs.

Server applications that can use per-processor licensing are treated differently. For example, when deploying SQL Server in virtual environments, customers can still utilize the multiple licensing models available in physical environments, such as per-processor or per-server/CALs. SQL Server Enterprise belongs to a select group of server applications, together with BizTalk Server Enterprise, Commerce Server Enterprise, and ISA Server Enterprise that are eligible for per-processor licensing. For these products, you can run unlimited virtual instances if licensed for all physical processors in the server.

Furthermore, per-processor licensing is based on the number of processors, regardless of how many cores are in the processor. For example, if you are using SQL Server Enterprise edition on a four-processor server with quad-core processors, you will require only four processor licenses. When virtualizing SQL Server, customers often consolidate to fewer physical servers or migrate to newer multicore physical servers and may end up requiring fewer processor licenses.

Applications that use the server/CAL model, such as Exchange Server, will require a server license for each running instance of the application. However, you may create and store an unlimited number of copies of the instance on any server or storage media. This can be a nonrunning copy, such as a backup or clone of a virtual machine.

TIP For more information about licensing and volume licensing Microsoft applications and operating systems, please refer to the Microsoft licensing website at `www.microsoft.com/licensing`. Also see "Application Server License Mobility" at `www.microsoft.com/licensing/resources/volbrief.mspx`.

Migrating to a Virtual Environment

The traditional method of deploying a new physical server generally takes weeks in many organizations. With VMware vSphere, you can deploy new virtual machines much more quickly in minutes from a template. For example, by creating an OS with all the patches and required updates and scripting SQL Server install, you can create a template that can be deployed over and over again, significantly reducing the time to provision SQL Server VMs.

When you create virtual machines, you can either use the VMware Converter tool or create fresh virtual machines leveraging vSphere's templates and clones.

Initial Creation of Virtual Machines

VMware Converter makes physical to virtual (P2V) migration very convenient and supports both hot and cold cloning of an existing physical server. You will need to watch out for two key issues with P2V, both of which can have a serious impact on VM performance:

Larger VMs Most server-class machines today have two sockets or more with at least a single-core processor in each socket. The newer machines have multicore processors. Add to this the fact that almost all physical servers are over-provisioned, with average CPU utilization at best in teens. Consequently, P2V invariably ends up creating larger VMs. Larger VMs means SMP-related virtualization overheads and reduced DRS effectiveness.

Extra processes within a VM Also, physical servers typically have administration tasks and services such as SNMP monitoring and backup agents running in background. These additional services can result in several performance issues such as IRQ conflicts and polling/interrupt-driven processes, all of which will typically exaggerate virtualization overhead after P2V.

When creating virtual machines using VMware Converter, you need to be careful about the previous two issues in P2Ved VMs. You should analyze and consider the validity of your physical server before using it as a candidate for a converter. Hence, the best-practice recommendation is to start with fresh VMs and take advantage of templates and cloning features in VMware vSphere.

In real life, you may not be able to avoid using VMware Converter in some cases. For example, you may have an old Windows server with some legacy OS/application stack, for which you cannot do a fresh install, recompile, or port to newer OS. Therefore, P2V is the best available option. In that case, watch out for the following:

- Before you power on the new VM after the P2V migration, edit the VM settings to reduce the CPU count and memory, adjust the Windows HAL to the UP version if needed, and remove any other peripherals not required.
- Before the P2V migration or after the P2Ved VM is powered on, turn off and disable all unneeded processes and services.

Ongoing Provisioning Strategy

Several server applications recommend scaling out horizontally using multiple virtual machines for scalability, load balancing, and high availability. In virtual deployment, you can power on or off virtual machines as needed to meet the usage demand. However, maintaining a large number of virtual machines can be challenging and can become confusing and time-consuming. For ongoing provisioning, you will need a strategy to create, maintain, and update the virtual machines and the server applications inside them. You take advantage of the templates and cloning features in VMware vSphere when provisioning multiple virtual machines. We discuss this in detail in Chapter 3.

Using templates, you can build a set of common server images with required Windows Server OS patches and updates. You can leverage VMware Update Manager to update your templates and VMs. We provide detailed steps to implement this later in this chapter. You can then customize these templates further by scripting a specific application installation and locking down their configurations as needed. You can easily clone or deploy these templates over and over again to reduce the time to provision additional VMs. You may need to create separate templates for different application roles and features, or if the servers are within a farm, you may need to save application-specific information locally. When you have multiple images (gold, silver, and so on), you will need a plan to deploy your templates efficiently.

Ensuring High Availability

When considering high availability for your virtualized applications or infrastructure, you must take into consideration the business requirements for application uptime, restoration, and continuance. Is automated failover, as in the case of an Exchange database availability group, required? Do you need to minimize downtime during patching? You should also review any existing service-level agreements that might be in place between you and your application owners. There are various means for providing high availability in the virtual environment, including the tools you are familiar with such as Microsoft Cluster Service and failover clusters and those unique to vSphere such as HA and Fault Tolerance. A truly high available solution will combine service availability with data availability and redundancy.

Most customers are shocked when they realize that failover clusters are not a true highly available solution. There is a single point of failure with storage. The storage itself won't fail, because all modern storage arrays are built with fully redundant subsystems that allow the array to continue to function in the event of component failure. However, there is the possibility of corruption through user and application inconsistencies. In stretched or geographically dispersed clusters with storage replication, corruption can be replicated to a remote data center. If procedures are not implemented to do data integrity checks, this corruption could go unnoticed for extended periods of time. When the need to fail over to your remote data center arises, you might find that because of the corruption nothing works.

Microsoft has implemented availability solutions that check data integrity and minimize the event of corruption due to application failure or inconsistencies. The newest of these technologies is *database availability groups* (DAGs), detailed in Chapter 6. DAGs use log shipping to transfer the latest transactions to the passive host where they are run through the integrity checking process and, assuming they are consistent, are then written to the database. In addition to this, there is the option to delay the commitment to the database by "lagging" the commitment. This is most useful in disaster recovery scenarios. In the following sections, we will look into each form of high availability that virtual machine and applications can take advantage of.

Host Availability

For those businesses whose minimal outages are due to server or host failure, there is VMware High Availability (HA). This feature provides out-of-the-box high availability for hosts and virtual machines where there is none. In the event of a host failure, all virtual machines on that host will be restarted on a surviving host, as shown in Figure 2.4. This is a perfect solution for those businesses that can withstand minor outages due to host failure. In this method of high availability, the application is not aware of the failure and will fail with the host.

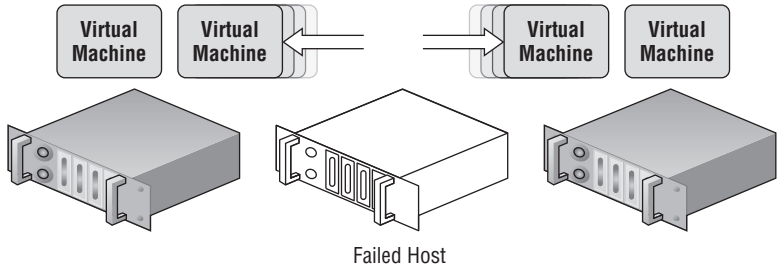

FIGURE 2.4 VMware High Availability

Application-Aware Availability

For those businesses that must maintain maximum application uptime, there are application-aware high-availability mechanisms such as the following:

- Failover clustering
- Database mirroring
- Log shipping for SQL Server and Exchange Server
- DAGs for Exchange 2010

Each of these application-aware mechanisms allows the application to detect failure within the operating system as well as the application and initiate the appropriate action to recover the application.

We will discuss failover clustering in Chapters 3 and 7. We also discuss Windows Network Load Balancing in Chapters 5, 8, and 9.

Host Availability Combined with Application-Aware Availability

In this hybrid model, VMware HA is protecting the hosts that the virtual machines running the application-aware availability solution are being hosted on.

In a simple scenario using database mirroring, the principal node is running at the primary data center and mirroring with a secondary node running at the remote data center or disaster recovery site.

The host with the principal database fails, and the mirroring fails over to the remote data center. In the time that it takes the mirroring session to fail over to the remote data center or disaster recovery site, VMware HA has restarted the virtual machine (formerly the principal node) on a surviving host at the primary data center without operator intervention.

The virtual machine is running and now available for replication and switchover. In this scenario, you have reduced the time that the database is operating in an unprotected state. Typically, in a physical server environment, the physical server would be unavailable until an operator diagnosed and restarted the operating system and SQL Server application.

This is just one scenario in which this hybrid solution can actually maximize your high-availability strategy.

Devising a Backup Strategy

Backing up Microsoft applications in virtual machines is very similar to performing backups in the physical environment. In some cases, backups performed in the virtual environment can be much simpler than in the physical environment. Typically, many customers will use the traditional method of deploying backup agents in the guest virtual machine and will perform streaming backup. There are numerous options for backing up your applications in the virtual environment. A properly designed and implemented backup and recovery plan can augment your high-availability and disaster recovery strategies.

A recommended best practice for backup tools in the virtual environment is to implement backup products that conform to the application programming interfaces (APIs).

Customers have deployed traditional application-level backup programs such as Symantec Backup Exec, HP Data Protector, and Commvault Galaxy products to conduct guest backups of their applications. These products are also capable of backing up the guest operating systems.

Using this method, you must fully be aware of the operating system running in the guest virtual machine and must employ the proper tools and recovery procedures to recover the virtual machine's operating system and applications that might be running in the virtual machine.

Application-level backups are well suited to backing up and recovering the following types of systems:

- Database systems
- Enterprise content management systems
- Custom applications
- Web servers
- Email and collaboration systems

Likewise, customers also deploy other third-party tools such as Vizioncore vRanger and vRanger Pro, VMware Data Recovery, Microsoft System Center Data Protection Manager, and other host-level backup products.

In the host-level backups method, you back up the virtual machines and all their encapsulated files. This backup type does not differentiate between the different types of operating systems running in the guest virtual machines.

Using this method, recovery is as simple as restoring the virtual machines to a functioning host and booting the virtual machines to get up and running again.

Backing Up Virtual Machines Using VMware Data Recovery

VMware Data Recovery (VDR) is a complete backup and recovery solution for executing host-level backups of your virtual infrastructure. VDR is available for those VMware customers who have purchased a vSphere Essentials Plus as well as Advanced, Enterprise, and Enterprise Plus licenses.

You will back up and restore virtual machines as a unit. This method will allow you to recover to a last point-in-time backup very quickly.

You will not worry with restoring the operating system and application. Depending on your business SLAs for Recovery Time Objective (RTO) and Recovery Point Objective (RPO), VDR can assist you in creating the perfect backup and recovery solution for your business organization.

In this section, you will do the following:

- Install the VMware Data Recovery virtual appliance
- Configure the VDR appliance
- Perform a backup of a Microsoft tier-1 application using VDR

Installing the VMware Data Recovery Appliance

Installing the VDR appliance is a straightforward process. You download the latest version of the VDR product from the VMware website. Unzip the downloaded file to your desktop or a server that you will perform the installation from.

On the Deploy OVF Template page, browse to the location on your desktop or server that you unzipped the installation package to, select the Deploy From File option, and browse to the OVF template file.

Once you have selected the OVF template file, click the Next button to continue the installation of VDR.

Figure 2.5 shows the vSphere vCenter, where you will select File ➢ Deploy OVF Template to install the OVF template file that will create the virtual appliance.

Once you have the OVF template file loaded, the installation will display the installation version and other details about the VDR version that you are installing. In Figure 2.6, you can see the details of the VDR version that you will be installing. The product information includes the following:

- The product name
- The version of the VDR product
- The vendor of the product
- The downloaded size of the template file
- The space that the appliance will consume when installed

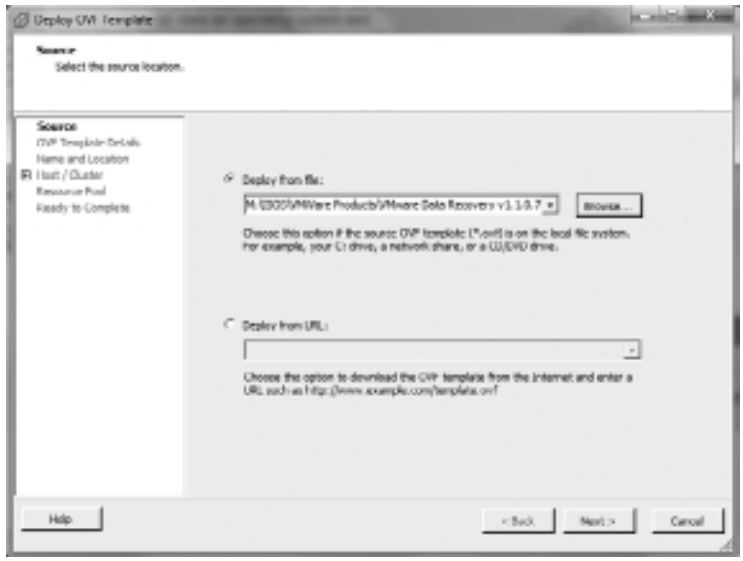

FIGURE 2.5
Installing the OVF template file

FIGURE 2.6
VDR virtual appliance installation information

Click the Next button to continue the installation. On the following screen, you will have the option of naming the VDR appliance and choosing where to locate the appliance in the vCenter inventory. After specifying your name and location, click the Next button to continue.

On the following screen, choose the host or cluster on which to run the VDR appliance. In Figure 2.7, you choose the cluster in which you will run the VDR appliance on. You can let vCenter choose the host with the most appropriate resources on which to run the VDR appliance.

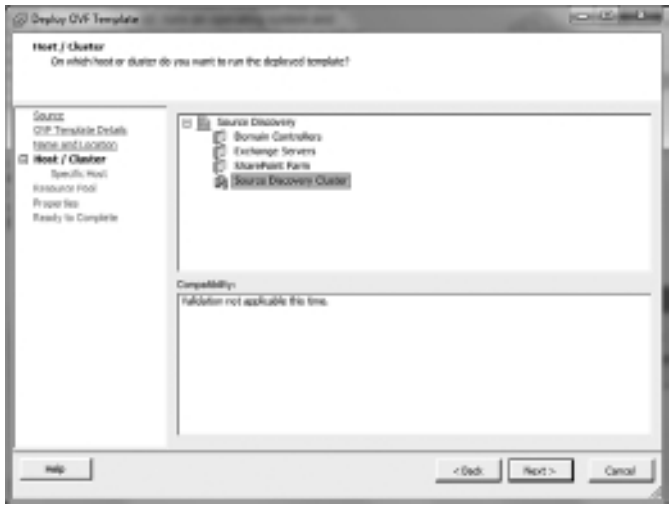

FIGURE 2.7
Choosing the cluster for the VDR appliance

Over the next screens, provide the following information that will be used for installing the VDR appliance:

- The host in the cluster for running the VDR appliance
- The IP address policy for the VDR appliance (fixed, transient, or DHCP)
- The time zone settings

Once you have specified all the required information, click the Next button, and you will see the Ready To Complete summary screen for your selections. This is your last chance to correct any of the information before you proceed with the installation.

Figure 2.8 shows the summary screen detailing your installation options. As you can see on the summary screen, we named our appliance VDR110707.

Once that's complete, you will proceed to configure the VDR appliance. In the next section, we will show how to configure the VDR appliance.

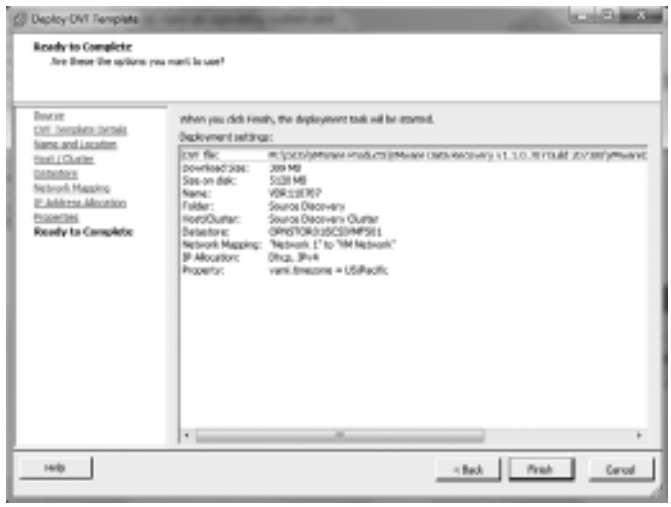

FIGURE 2.8
VDR Ready To Complete summary screen

Configuring the VMware Data Recovery Appliance

In this section, you will complete the following configuration tasks to prepare the VDR appliance for your first tier-1 application backup:

1. Power on the VDR appliance for the first time.
2. Connect to the VDR appliance for the first time.
3. Add a backup destination.

To start the configuration of the VDR appliance, first power on the appliance. From the vSphere vCenter inventory, right-click the VDR appliance, select Power, and then select the Power On context option to power on the VDR appliance. If you want to log into the appliance, then you can open the appliance by right-clicking the appliance and selecting the Open Console context option. Figure 2.9 shows powering on the VDR appliance.

Figure 2.9
Powering on the VDR appliance

What is a Virtual Machine?
Basic Tasks
Power on the virtual machine
Edit virtual machine settings
Explore Further
Learn more about virtual machines
Learn how to install an operating system

Figure 2.10 shows the console of the VDR appliance. Here you can log into the appliance and configure the appliance from here.

You will get such information as the URL from which to administer the appliance and other options for configuring the VDR appliance. We will show how to configure the VDR appliance from the vCenter.

You can now move back to the vCenter to continue configuring the VDR virtual appliance. During the VDR appliance installation, a plug-in to manage the appliance is installed on the workstation or server that you performed the installation from as well as into the vCenter. The first thing you will need to verify is whether the VDR plug-in is installed and enabled in the vCenter.

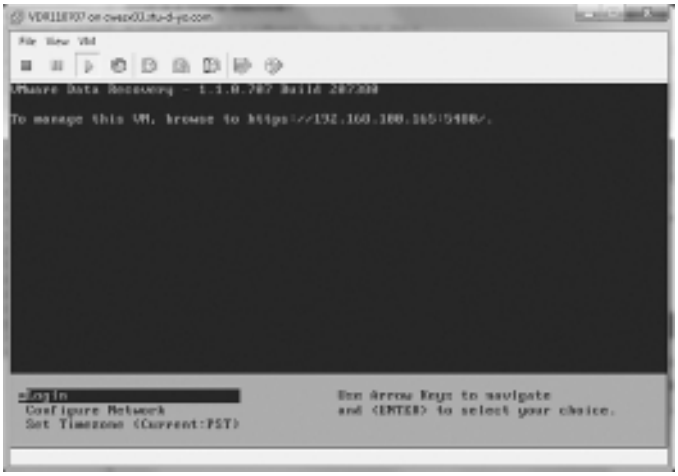
VMware Data Recovery - 1.1.0.787 Build 287388
To manage this VM, browse to https://192.168.100.165:5480/.
Login
Configure Network
Set Timezone (Current:PST)
Use Arrow Keys to navigate
and <ENTER> to select your choice.

FIGURE 2.10
Examining the VDR console

From the vCenter options bar, select the Plug-ins options, and click the Manage Plug-ins context option. This will display the status of the installed plug-ins that your vCenter knows about. In Figure 2.11, you can see the VDR plug-in and its status shown as Enabled.

FIGURE 2.11
VDR managed plug-in status

If the VDR plug-in is not installed automatically, you will have to manually install it from the location that you unzipped the downloaded file. The installation is straightforward. Double-click the installer file, and follow the instructions for installing the VDR plug-in. You will need to exit the vSphere vCenter after the installation has completed to view the plug-in.

Now that you know you can manage the VDR appliance, you can now start to configure the VDR appliance from the vCenter. From here, you will now connect to the VDR appliance to start the configuration. To connect to the appliance, click the Inventory option in the navigation bar. This will bring four options on the context menu. Select the Solutions And Applications option, and then select the VMware Data Recovery option. Figure 2.12 displays the procedure for connecting to the VDR appliance.

When you open the VMware Data Recovery application for the first time, you will be presented with the welcome screen. It is from this screen you will select and connect to the VDR appliance for the first time.

You will see that no appliance is selected to connect to on the welcome screen. In the inventory list of virtual machines, select the VDR appliance, and the empty box will be filled in with the VDR appliance name.

Figure 2.13 displays the name of the VDR appliance populated in the Name, IP Address, or DNS field. At this point, you will click the Connect button to connect to the VDR appliance for the first time.

FIGURE 2.12
Connecting to the VDR application

FIGURE 2.13
Connecting to VDR appliance

TIP If you cannot connect to the VDR appliance, try resetting the appliance by right-clicking the appliance in the inventory and choosing Reset from the context menu.

Upon connecting to the VDR appliance for the first time, you will be asked to verify the credentials of the administrator of the appliance to the vCenter. Type the password for the vCenter.

Figure 2.14 displays the pop-up you will encounter when first connecting to the VDR appliance.

FIGURE 2.14
Confirming vCenter credentials

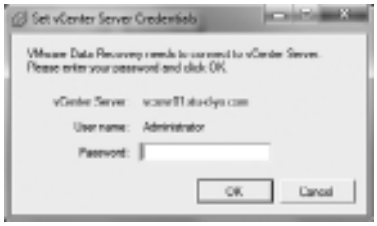

After confirming the credentials of the vCenter, the VDR Getting Started Wizard will start, and you will again be asked for your vCenter credentials. Enter the same password, and the wizard will continue. Figure 2.15 shows the VDR Getting Started Wizard where you will again enter your vCenter credentials to continue.

FIGURE 2.15
The Getting Started Wizard

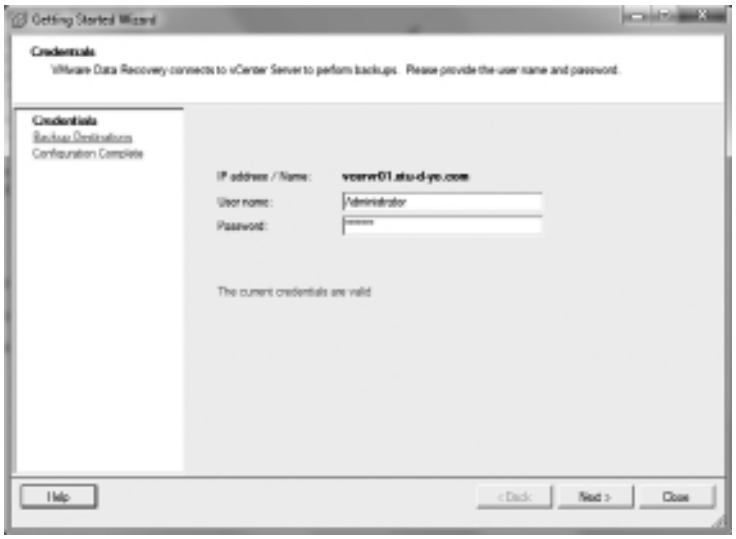

The next screen will allow you to select a backup destination. A backup destination can be a network file share, a raw device mapping (RDM), or a VMFS data store. You can also add, format, and remove backup destinations from this screen.

The backup destination is where all your virtual machine backups will be stored and will also be the source of deduplication. Figure 2.16 shows the backup destination screen. There will be a little more configuration for you to do if you have opted to use an RDM for your backup destination.

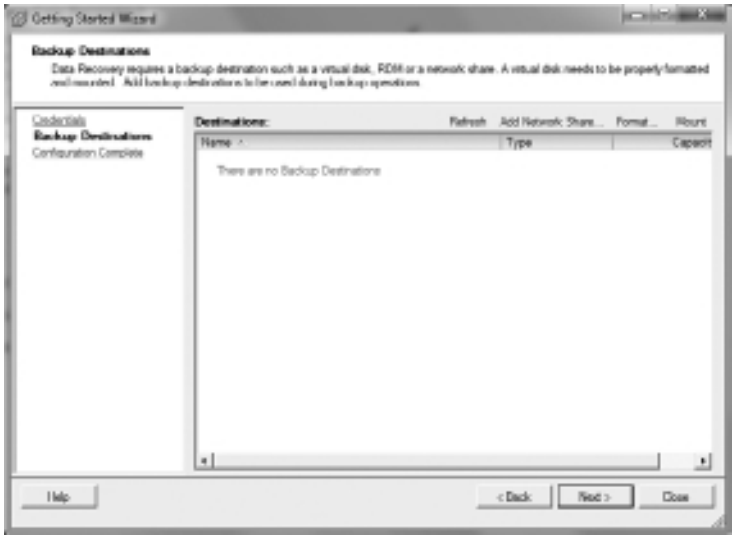

FIGURE 2.16 Configuring backup destinations

Clicking the Next button will take you to the Configuration Complete screen, as shown in Figure 2.17, at which point you can click the Close button to finish the configuration.

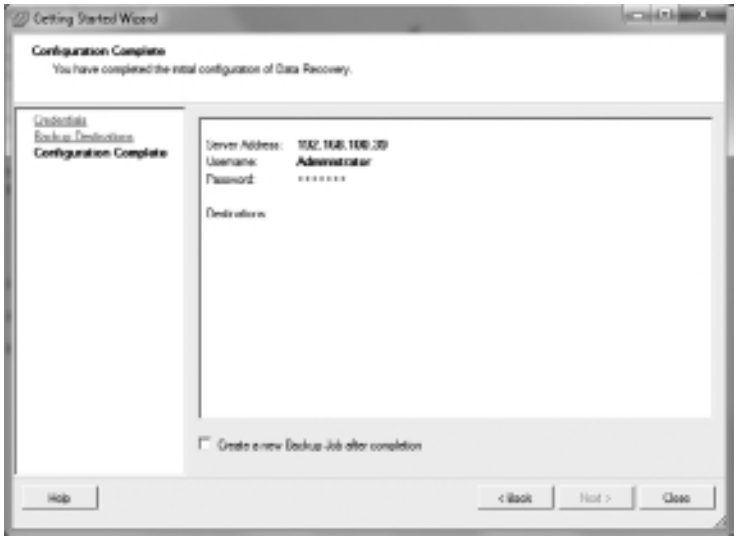

FIGURE 2.17 Finishing the configuration

Because you have opted to use RDMs in this example, you must do a little more configuration. The first thing you must do is to edit the virtual machine configuration to add the RDM to the virtual appliance's storage. You do this by exiting the VDR application by selecting the Solutions And Applications option on the navigation bar, choosing the Inventory option, and selecting the Hosts And Clusters option.

Once in the vCenter inventory, select the virtual appliance, and under the Basic Tasks area, click Edit Virtual Machine Settings. Here you will be able to edit the virtual machine's configuration. Figure 2.18 shows the VDR appliance virtual machine properties.

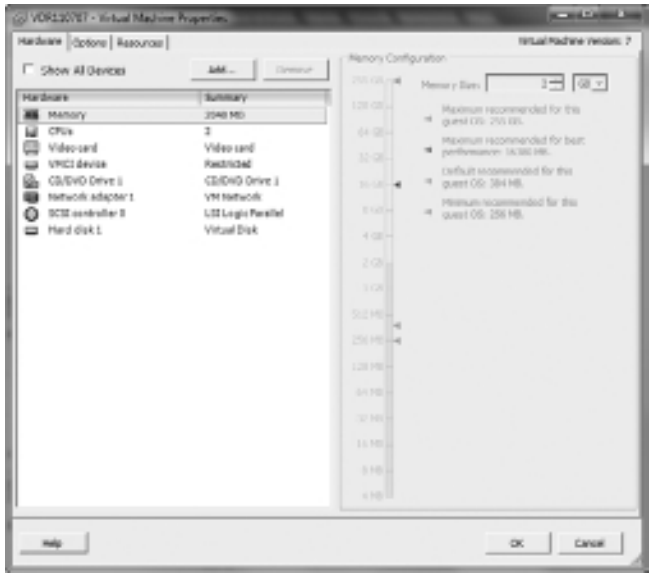

FIGURE 2.18
VDR Virtual Machine Properties page

From the Virtual Machine Properties page, you can add a hard disk to your configuration. Instead of choosing a VMFS data store, select the Raw Device Mappings (RDM) option, and select the storage to present as an RDM to the virtual machine. Figure 2.19 shows the completed configuration after the RDM has been added to the virtual machine configuration.

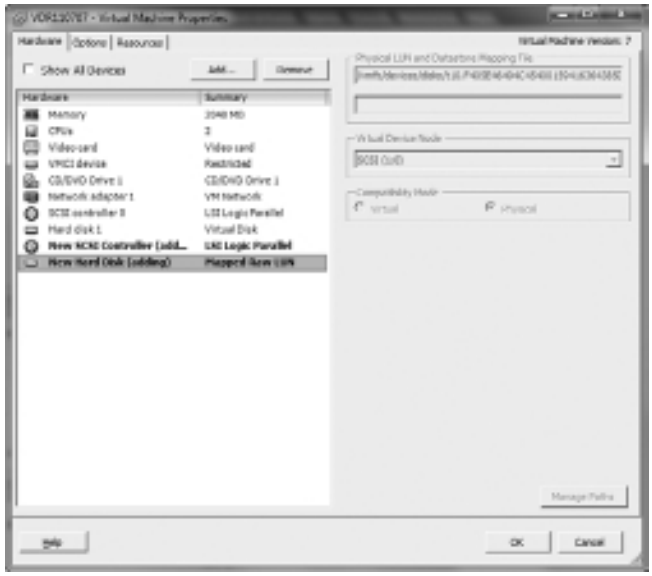

FIGURE 2.19
Completed virtual machine configuration

NOTE The configuration of RDMs and other types of storage will be covered in Chapter 3. Please refer to that chapter for additional information if needed.

Now that you have mapped the RDM to the VDR virtual machine, you can reconnect to the VDR application and configure a new backup destination using the RDM that you just added to the virtual machine.

Reconnect to the VDR application, as shown earlier in Figure 2.13. After reconnecting to the VDR application, select the Configuration tab and the Destinations option to the left of the menu. Once you select the Destinations option, you will see the RDM you added to the VDR appliance.

At this point, you will select the new backup destination and choose the Format option. This will start the format operation in the vCenter. Figure 2.20 displays the format backup destination warning.

Format
Format this volume to be used as a Backup Store. Click OK to format and partition the volume as a backup destination.
OK
Cancel
Help

FIGURE 2.20 Format backup destination warning

WARNING To prevent data loss, ensure that the chosen backup destination does not contain any data that might be considered as valuable to your business operations. The format operation will format the chosen backup destination, destroying all data that was previously on disk.

After you acknowledge the format warning, as shown in Figure 2.21, the format of the backup destination will proceed. When complete, you will now have a single backup destination to start storing your virtual machine backups to.

FIGURE 2.21 Formatting the backup destination disk

Figure 2.22 shows a completed formatted backup destination. You can now proceed to back up your tier-1 application.

FIGURE 2.22 Completed format operations on the backup destination

PERFORMING A BACKUP OF A TIER-1 APPLICATION

You can now select the Backup tab to create a backup job. Once in the backup area, you select the New option to create a new backup job. You will be presented with an inventory listing of all your virtual machines. Select which virtual machines to back up.

WARNING When selecting virtual machines to back up, be careful not to select the Active Directory domain controllers. VDR uses VSS-enabled snapshots to backup entire virtual machines. Snapshots of Active Directory domain controllers can cause corruption. Refer to Chapter 4, "Virtualizing Active Directory," for more information on backup procedures for Active Directory domain controllers.

Figure 2.23 shows the virtual machine inventory that you can choose from when selecting the virtual machines to back up in your backup jobs. In this example, we have selected five virtual machines to back up.

After you choose your virtual machines to include in your backup job, you must then select a backup destination for your backups. In this case, this would be the RDM that we added to the virtual appliance.

From there you can move to the selection of the backup window. This determines when to perform backups. From there, you move to the backup retention policy page in which you can select how long and how many copies of the backups you keep.

After you configure the retention policy, you now come to the Ready To Complete screen where you will save the backup job. Mind you, you have not started the job yet. Figure 2.24 shows the summary page.

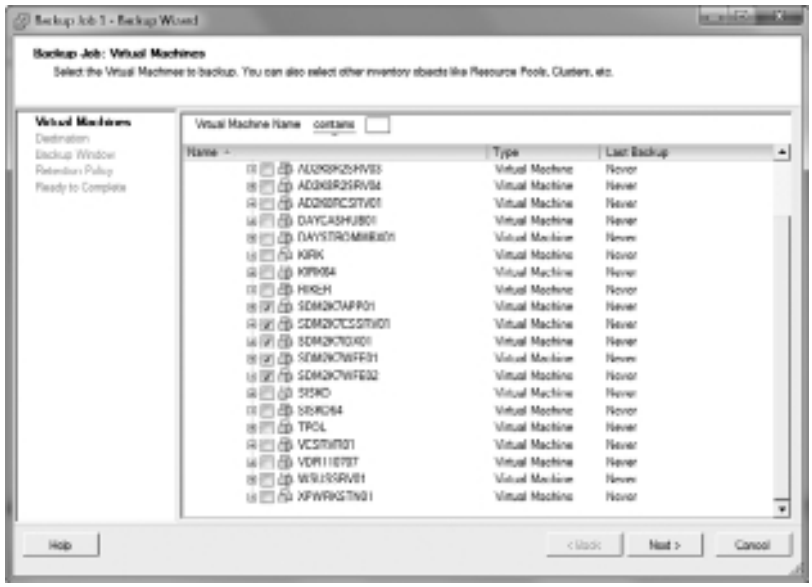

FIGURE 2.23
Virtual machine inventory

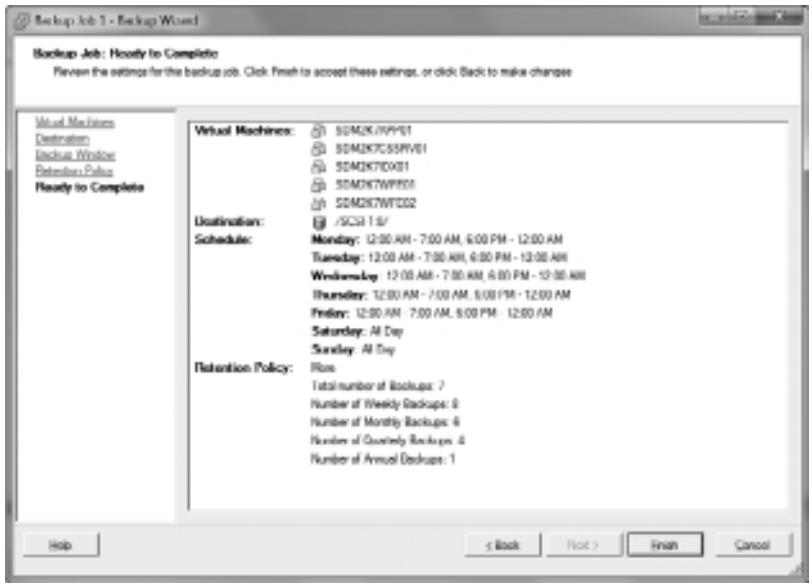
Backup Job: Ready to Complete
Review the settings for this backup job. Click Finish to accept these settings, or click Back to make changes
Virtual Machines
Destination
Backup Window
Retention Policy
Ready to Complete
Virtual Machines:
Destination:
Schedule:
Monday: 12:00 AM - 7:00 AM, 6:00 PM - 12:00 AM
Tuesday: 12:00 AM - 7:00 AM, 6:00 PM - 12:00 AM
Friday: 12:00 AM - 7:00 AM, 6:00 PM - 12:00 AM
Saturday: All Day
Sunday: All Day
Retention Policy:
Total number of Backups: 7
Number of Weekly Backups: 6
Number of Monthly Backups: 6
Number of Annual Backups: 1
Help
< Back
Next >
Finish
Cancel

FIGURE 2.24
Saving the created backup job

After the backup job configuration is complete, you are now ready to start a backup. Select the Backup tab, and find your saved backup job. After your backup job has been saved, you will notice that the virtual machines that were selected for backup are not in compliance. What does this mean?

This simply means that you have never performed a backup on these virtual machines. To start your backup, you will select the saved backup job, right-click, and select the Bring Into Compliance option from the context menu. This will start the backup job. Figure 2.25 displays the context menus.

FIGURE 2.25
Starting the backup job

Now that the backup job is running, you can view the active status of the running backup job by selecting the Reports tab and, under the Report Name field, selecting the Running Tasks option. Figure 2.26 displays the backup job's running status.

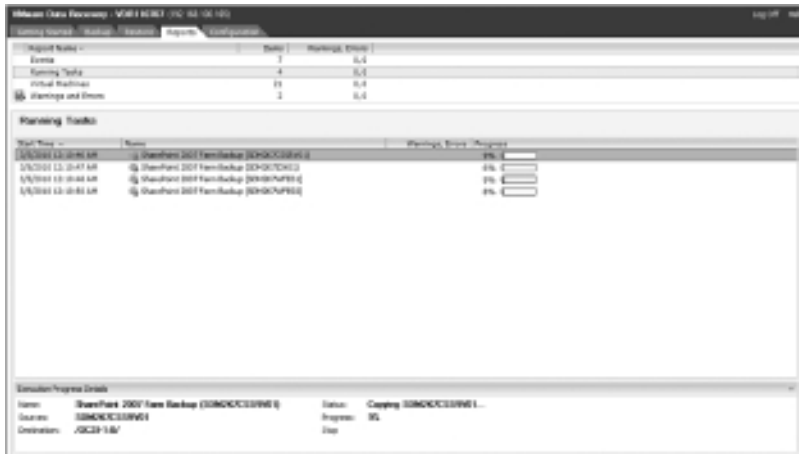

FIGURE 2.26
Monitoring the backup job's status

When the backup has been completed, you can check to ensure that your virtual machines were backed up successfully by selecting the Reports tab. Under the Report Name field, select the Events option. Figure 2.27 displays the backup job's completion status.

FIGURE 2.27
Displaying the backup job's completion status

Patching Virtual Machines Using VMware Update Manager

VMware provides a patch management framework for VMware vSphere through its Update Manager product. You can use Update Manager to do the following:

- Patch and upgrade VMware ESX/ESXi hosts
- Apply patches to Windows and certain versions of Linux in virtual machines
- Upgrade VMware Tools and virtual hardware for virtual machines
- Patch and upgrade virtual appliances

VMware Update Manager integrates with the Shavlik patch management technology to provide these patching capabilities. You can also patch offline virtual machines by using a special patching network for suspended or offline systems. In vSphere deployments, you can leverage the Update Manager option instead of Microsoft Software Update Services (SUS) or its successor, Windows Server Update Services (WSUS) for your patching and update requirements.

Update Manager also allows you to create update baselines that represent a set of updates and then compare a host or VM against that baseline. If the host or VM is not compliant against that baseline, then it can be remediated to deploy the updates. This can be very useful to ensure consistency across several Windows virtual machines in your vSphere environment.

Deploying VMware Update Manager

You can deploy Update Manager in three different ways based on the size of your vSphere deployment:

- vCenter Server and Update Manager servers as well as their database instances are installed on the same host. This model requires only one host and is recommended for relatively small virtual environments.
- vCenter Server and Update Manager servers are installed on one host, but their database instances are deployed on two separate hosts. This model needs three hosts and is recommended for medium deployments, with more than 300 virtual machines or 30 ESX hosts.

- vCenter Server and Update Manager servers run on different hosts, and each one has a dedicated host for its own database instance. This model requires four hosts and is recommended for large deployments when the data centers contain more than 1,000 virtual machines or 100 hosts.

Because patching can be disk I/O intensive, you should provide the Update Manager server host with enough memory to cache patch files within the system cache. VMware recommends using at least 2 GB of RAM for the Update Manager server host. Another recommendation is to separate the patch store and Update Manager database on separate physical disks to improve the Update Manager I/O performance.

VMware Update Manager Sizing Estimator

VMware provides a sizing tool, VMware vCenter Update Manager Sizing Estimator, to estimate Update Manager database size, patch store disk space, and temporary disk space. The tool is available at `www.vmware.com/support/vsphere4/doc/vsp_vum_40_sizing_estimator.xls`. The tool provides recommendations about Update Manager server and database deployment models, initial disk usage for database, patch store, and temporary space and monthly disk usage growth for database and patch store.

You can also deploy vCenter and Update Manager (Figure 2.28) in virtual machines to leverage virtualization benefits. For the proof-of-concept deployment for Source Discovery Corporation, which we discuss in detail in later chapters, you'll use the single host deployment model mentioned here.

Figure 2.28
VMware vCenter Update Manager

Using VMware Update Manager

Using VMware Update Manager to upgrade and apply patches is a multistep process. You need to follow these steps to ensure a smooth update with a minimum of system downtime:

1. Configure the Update Manager patch download source.
2. Download patches, extensions, and related metadata.
3. Create baselines and baseline groups.
4. Attach baselines and baseline groups to vSphere objects.
5. Scan selected vSphere objects.
6. Remediate selected vSphere objects.

In the following sections, we will detail the process to patch the Windows virtual machines in the Source Discovery proof of concept.

Configuring the Update Manager Patch Download Source

The Update Manager process begins by downloading information about a set of patches and extensions. So, the first step is to configure the Update Manager server to download patches and extensions either from the Internet or from a shared repository, as shown in Figure 2.29. You can do this from the Patch Download Settings on the Configuration tab. Figure 2.29 shows that we configured Update Manager to download only Windows patches from Shavlik.

Figure 2.29
Configuring patch download source

Downloading Patches, Extensions, and Related Metadata

Next you need to download the patches, extensions, and related metadata. You can either do this manually or automate it so that Update Manager contacts Shavlik, VMware, or third-party sources at regular configurable intervals to gather the latest information (metadata) about available patches or extensions. Use the Schedule Update Download Wizard (Figure 2.30) under Patch Download Settings to configure this. For the proof of concept, we've set the patch downloads to occur at noon every day. You can change to this a weekly or higher interval per your requirements.

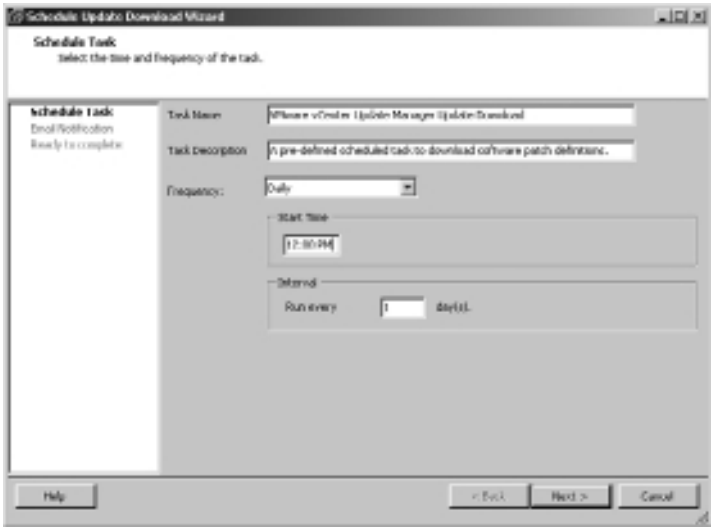

FIGURE 2.30
Patch download schedule

Creating Baselines and Baseline Groups

Update Manager allows you to combine one or more patches, extensions, service packs, bug fixes, or upgrades to form a baseline. Baselines make it convenient to apply multiple updates at the same time. You can also combine multiple baselines into a baseline group. You can use baseline groups to scan and remediate a vSphere object against all of them as a whole. Creating baselines and baseline groups is an optional step but is highly recommended for easier administration.

For the Source Discovery proof of concept, you first start by creating a baseline that includes all Windows Server 2008 patches. From the Baseline and Groups tab, you click the Create link next to the Patch Baselines tab to start the baseline creation wizard, as shown in Figure 2.31. Enter **Windows Server 2008 Baseline** as the baseline name, and select VM Patch as the baseline type. VM patch baselines allow you to patch virtual machines based on selected patch criteria.

You leave the baseline type as Dynamic, as shown in Figure 2.32. Dynamic baselines are updated automatically if new patches that meet the baseline criteria are added to the patch database.

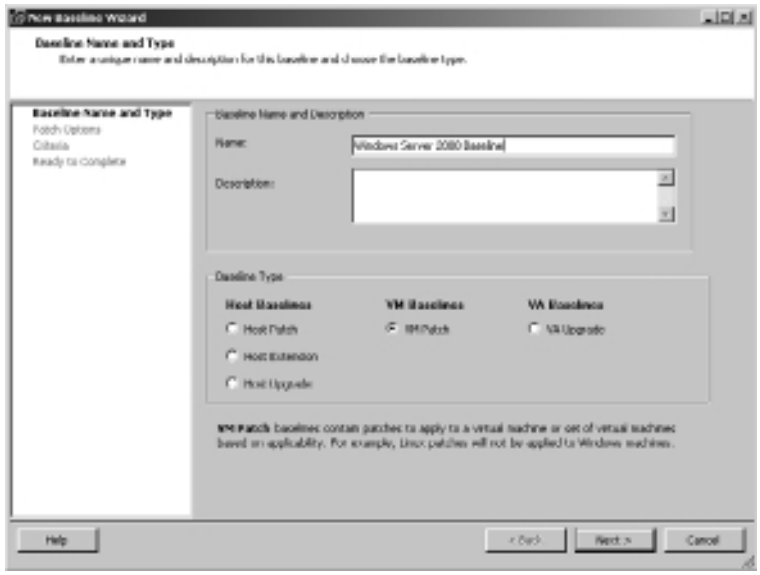

FIGURE 2.31
Creating a patch baseline

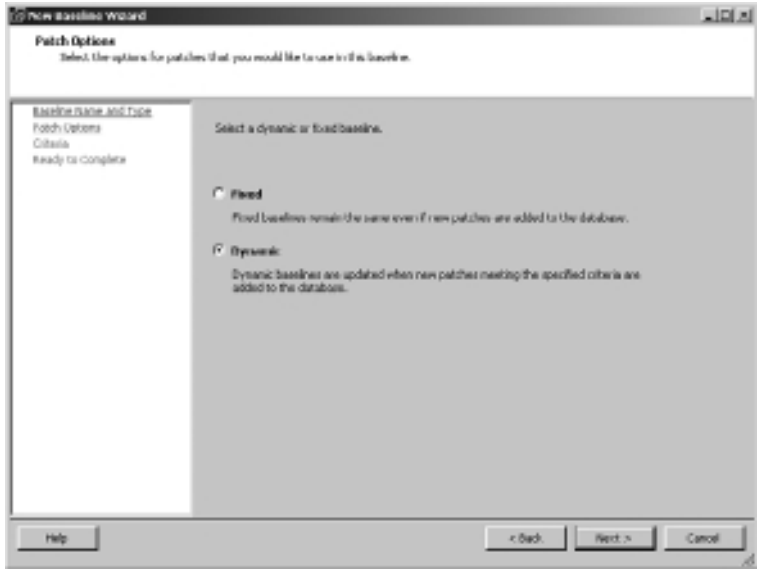

FIGURE 2.32
Baseline type

The Source Discovery proof of concept uses Windows Server 2008 as a guest operating system. So, you search for all products containing a *Windows Server 2008* text string, as shown in Figure 2.33. You can select multiple products by pressing the Ctrl key simultaneously with the product names.

FIGURE 2.33
Baseline criteria

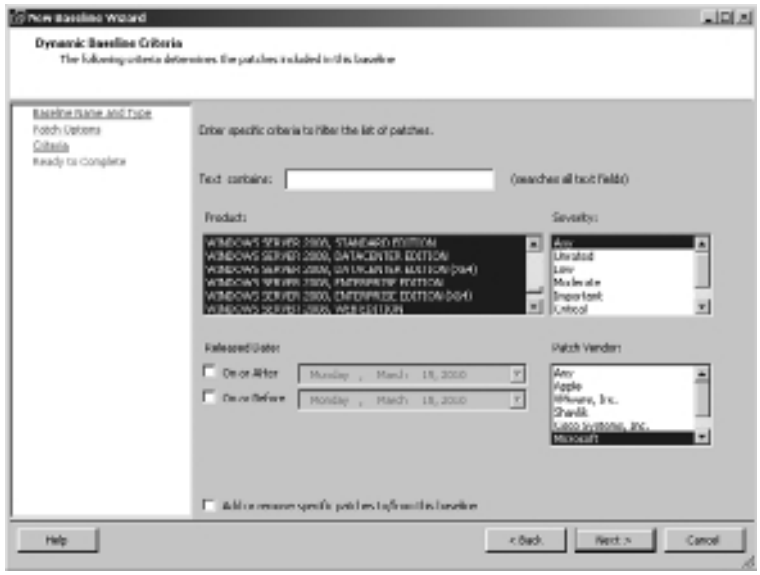

On the next screen, shown in Figure 2.34, you get to review your selections. Click Finish to create the baseline.

FIGURE 2.34
Baseline review selections

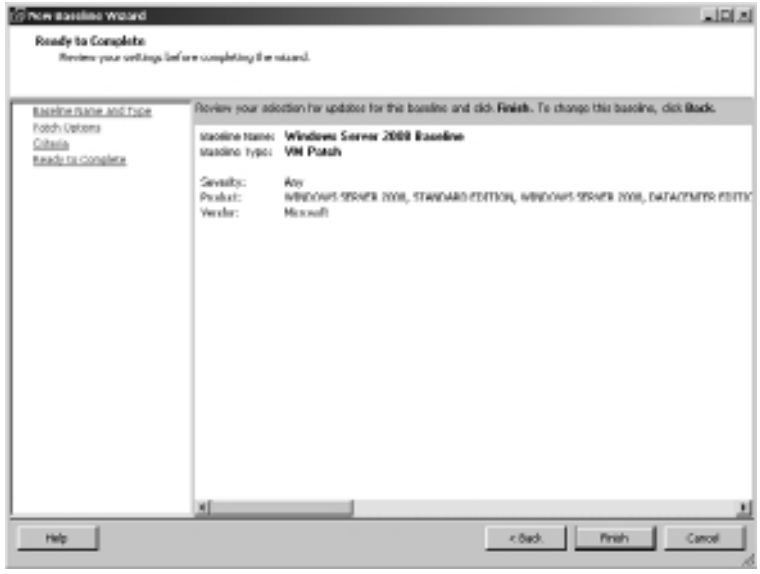

The Update Manager screen shows the newly created Windows Server 2008 baseline, as shown in Figure 2.35.

FIGURE 2.35
Update Manager VM patch baselines

Next you can create a baseline group to help you apply Windows Server 2008 patches, upgrade VMware tools for the proof-of-concept virtual machines, and upgrade the VM hardware version with a single task. This is particularly convenient when you have to patch multiple virtual machines in a data center. To create a baseline group, you click the Create link next to Baseline Groups on the Baseline And Groups tab. In the baseline group creation wizard (Figure 2.36), you enter **Windows Server 2008 Baseline Group** as the baseline group name and select Virtual Machines And Virtual Appliances Baseline Group as the baseline group type.

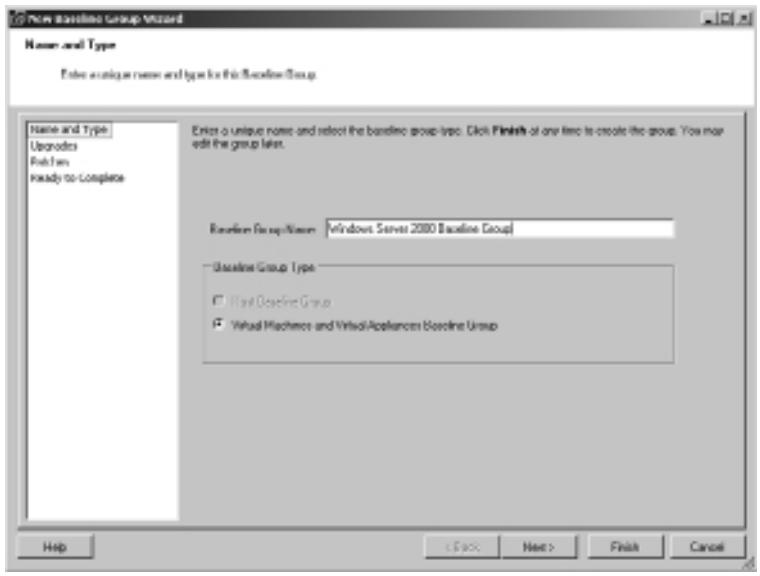

FIGURE 2.36
Creating a baseline group

You select the VM hardware and VM Tools upgrades. If a baseline group contains both upgrade and patch baselines, the upgrade runs first. Update Manager will power off the virtual machine before upgrading the virtual hardware, so upgrading virtual machine hardware is faster if the virtual machine is already powered off. Similarly, Update Manager will power on the virtual machine before the VMware Tools upgrade. So, upgrading VMware Tools is faster if the virtual machine is already powered on. Select VMware Tools Upgrade To Match Host (Figure 2.37).

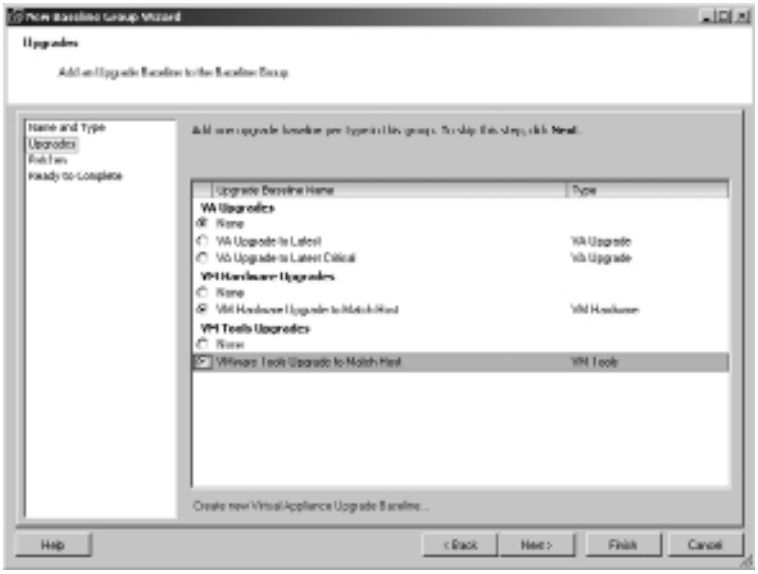

FIGURE 2.37
Adding upgrade baselines to a baseline group

You then add the Windows Server 2008 patch baseline to the baseline group, as shown in Figure 2.38. Notice that you can add multiple patch baselines to a single baseline group.

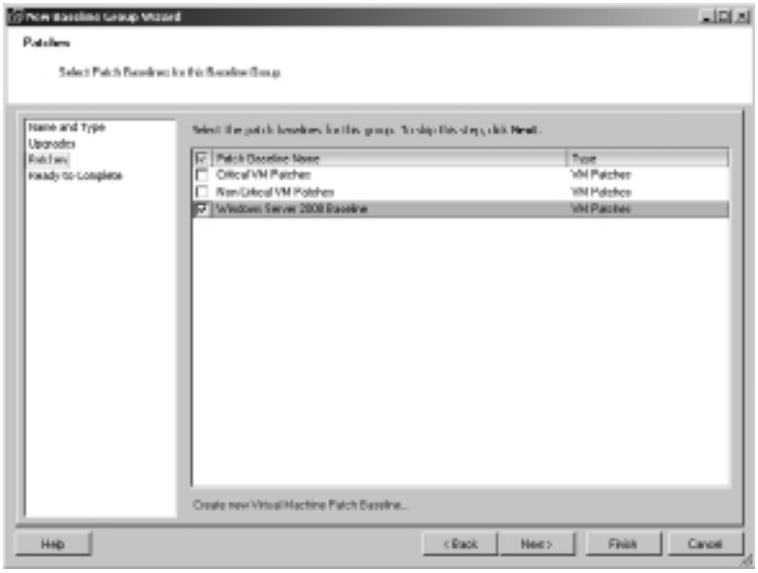

FIGURE 2.38
Adding patch baselines to a baseline group

On the next screen, shown in Figure 2.39, you get to review your selections. Click Finish to create the baseline group.

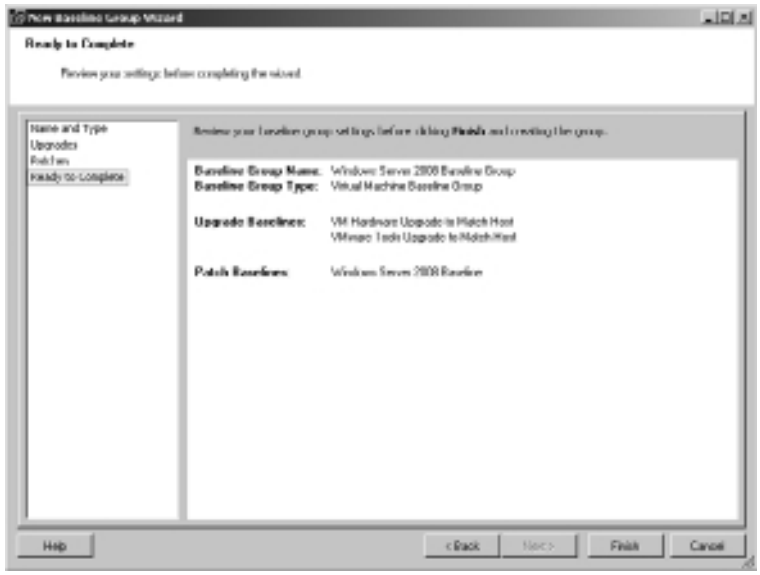
New Baseline Group Wizard
Ready to Complete
Review your settings before completing the wizard
Name and Type
Upgrades
Patches
Ready to Complete
Review your baseline group settings before clicking Finish and creating the group.
Baseline Group Name: Windows Server 2008 Baseline Group
Baseline Group Type: Virtual Machine Baseline Group
Upgrade Baselines: VM Hardware Upgrade to Match Host
VMware Tools Upgrade to Match Host
Patch Baselines: Windows Server 2008 Baseline
Help
< Back
Next >
Finish
Cancel

FIGURE 2.39
Baseline group review selections

The Update Manager screen shows the newly created Windows Server 2008 baseline group, as shown in Figure 2.40.

FIGURE 2.40
Update Manager baseline groups

Attaching Baselines and Baseline Groups to vSphere Objects

To use baselines and baseline groups, you must attach them to specific vSphere objects such as ESX hosts, virtual machines, or virtual appliances. After you select an object from the left inventory pane, you click the Attach link on the Update Manager tab. In Figure 2.41, you can see that you can attach the Windows Server 2008 baseline group created earlier.

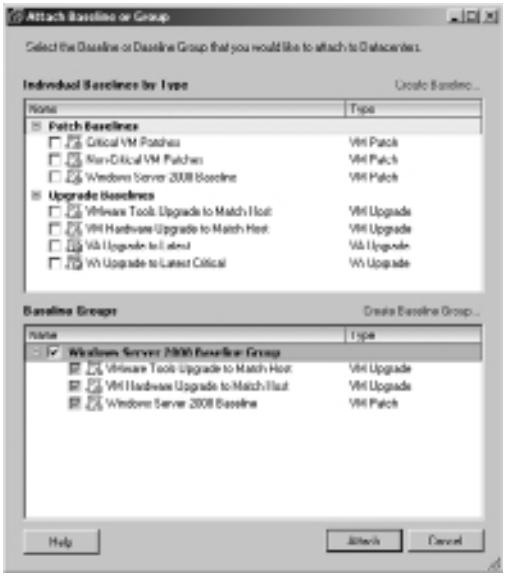

FIGURE 2.41
Attaching a baseline or baseline groups

Scanning Selected vSphere Objects

You can scan a set of vSphere hosts, virtual machines, or virtual appliances to determine how they comply with baselines and baseline groups that you attach. The scan process can be started manually or through scheduled tasks. You right-click a vSphere object in the left inventory pane and select the Scan For Updates menu option, as shown in Figure 2.42. Alternatively, you can select the object and click Scan link on the Update Manager tab to verify compliance.

FIGURE 2.42
Scanning a virtual machine for updates

When a powered-on scan is run on a Windows virtual machine for the first time, the Update Manager guest agent is installed in the Windows virtual machine. Therefore, the first powered-on scan can take longer than subsequent scans, and you should plan for this additional time when you run the first scan command. Once the scan process is complete, the Compliance view shows the attached baselines or baseline groups and compliance status based on them (Figure 2.43).

FIGURE 2.43
Compliance view

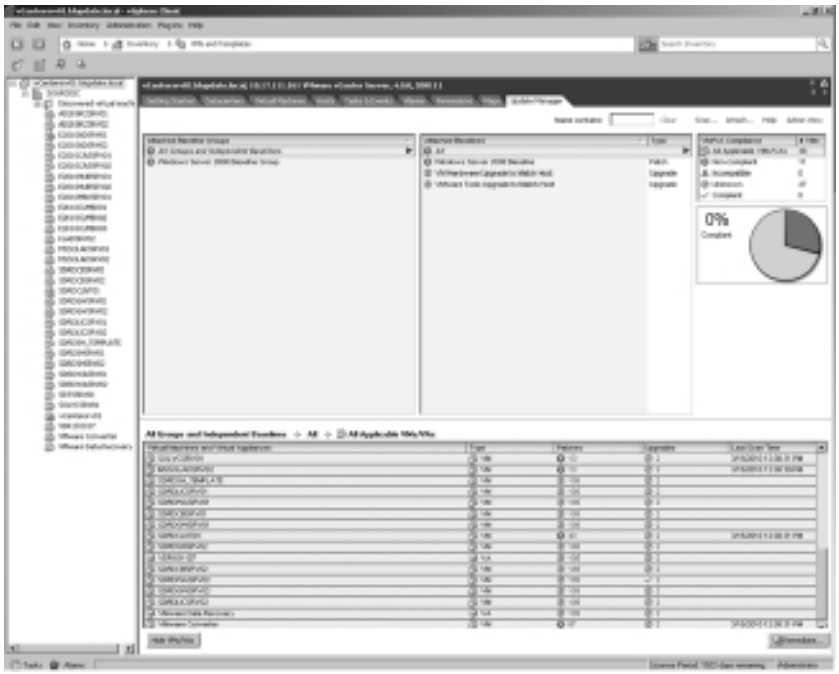

The compliance view responsiveness will depend upon the number of attached baselines. For large vSphere deployments, you should remove unused baselines to avoid a slow response.

REMEDIATING SELECTED VSPHERE OBJECTS

Update Manager can remediate vSphere objects by applying patches, extensions, and upgrades defined in the attached baseline or baseline groups to ensure that these objects are compliant. Remediation can help you keep vSphere hosts, virtual machines, and appliances up-to-date with required fixes and protect your virtual deployment against known security attacks. To remediate an object, you right-click the object in the left inventory pane and select the Remediate menu option. Alternatively, you can click the Remediate button in the Update Manager Compliance view for the object.

The Remediate Wizard allows you to select specific baselines or baseline groups and target objects for remediation, as shown in Figure 2.44. You select the Windows virtual machines used in the Source Discovery proof of concept as target objects, and you use the Windows Server 2008 baseline group for remediation.

FIGURE 2.44
Remediation target objects selection

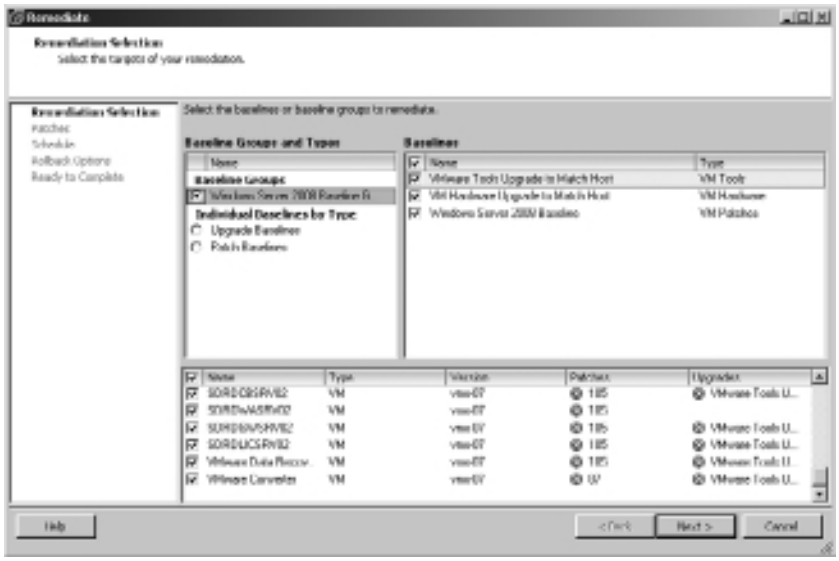

Next, the wizard allows you to select specific patches that you want to apply, as shown in Figure 2.45. For dynamic baselines, the list of patches displayed may change depending upon when the actual remediation occurs.

FIGURE 2.45
Remediation patches selection

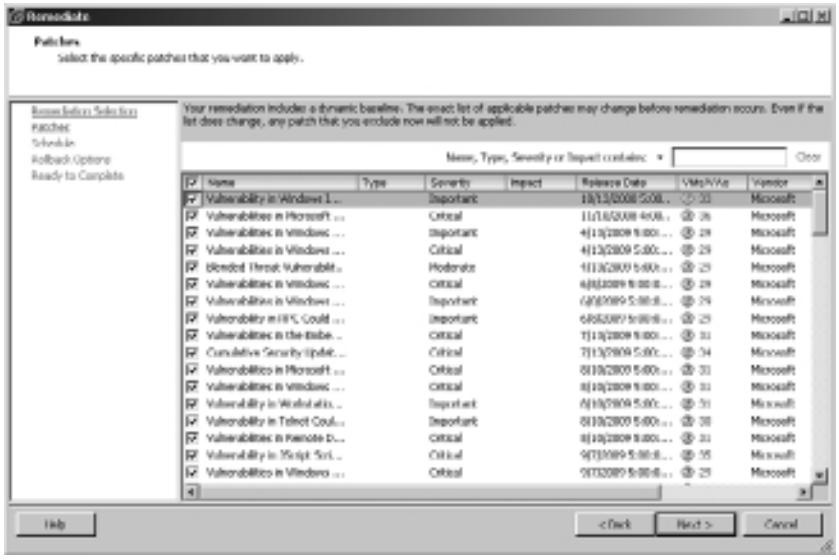

You can remediate the target objects immediately or schedule the remediation process, as shown in Figure 2.46. The virtual machines need to be powered on for applying patches and

upgrading VMware Tools. For a VM hardware upgrade, the virtual machines need to be powered off.

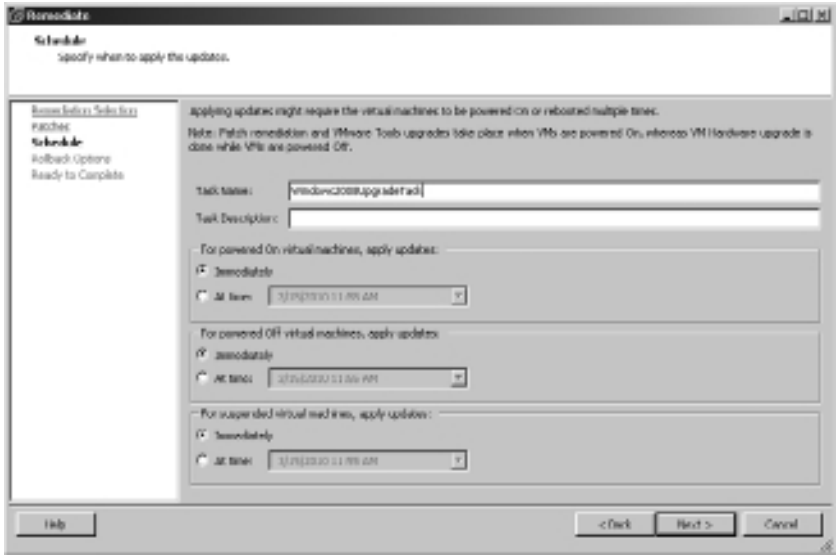

FIGURE 2.46 Remediation schedule

The Remediate Wizard allows you to snapshot virtual machines before remediation so that you can roll back any updates if needed, as shown in Figure 2.47. If enabled, the rollback option provides you with an option to keep the snapshots around for additional time until you are comfortable with the remediation.

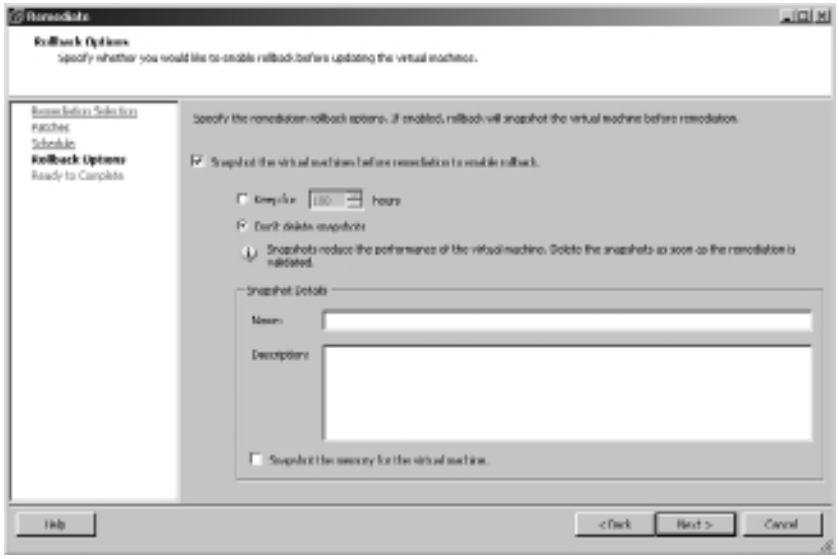

FIGURE 2.47 Remediation rollback options

You can review your remediation (Figure 2.48) options and click Finish to start or schedule remediation process.

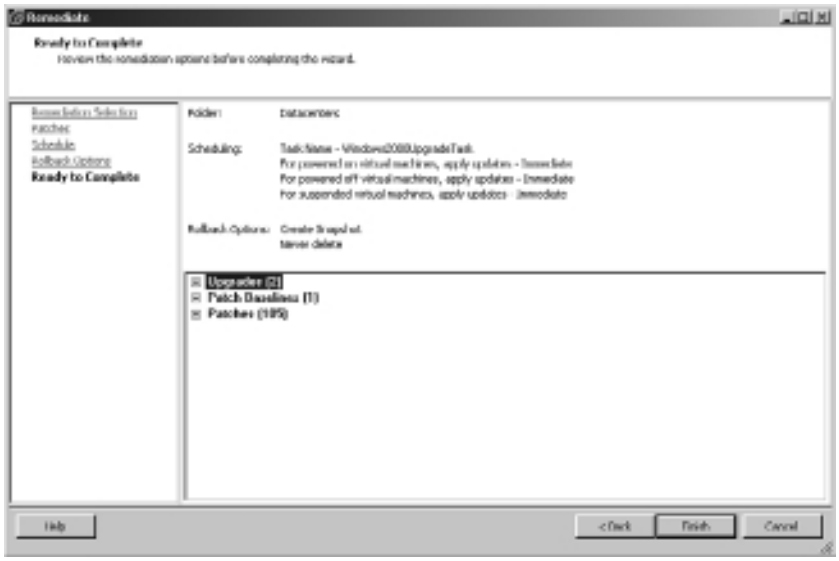

FIGURE 2.48 Remediation review options

Microsoft Server Virtualization Validation Program

In August 2008, Microsoft launched the Microsoft Server Virtualization Validation Program (SVVP) to validate third-party hypervisors for running Microsoft products. Hypervisors that pass validation under this program receive official support from Microsoft to run the Windows Server operating systems as well as several other server applications. To ensure that the configuration is supportable by Microsoft Support, Microsoft requires the virtualization solution to be tested on server hardware that is certified for Windows Server.

The supported applications include Windows Server and the most commonly deployed tier-1 applications such as the following:

- Active Directory 2000 (x86), 2008, R2
- SQL Server 2005, 2008, R2
- Exchange Server 2003, 2010
- SharePoint Server 2003, 2007
- BizTalk Server 2004, 2009
- ISA Server 2004, 2006
- Dynamics AX, CRM v4.0, NAV 2009

VMware ESX was the first hypervisor to pass validation under the Microsoft SVVP program. Starting from VMware ESX 3.5 Update 2, VMware and Microsoft are collaborating to expand support to other ESX versions. Currently, VMware ESX/ESXi 3.5 Updates 2 through 5 and vSphere 4.0 and 4.0 Update 1 are validated under the SVVP program. Customers can now receive full collaborative support from Microsoft and VMware when running their Microsoft applications on ESX.

For more information on the Microsoft Server Virtualization Validation Program, see the Microsoft Windows Server Catalog website at `www.windowsservercatalog.com`. Click the SVVP button in the top-right corner of the web page, and proceed to the Support Policy Wizard. In the next few paragraphs, we will walk you through the wizard steps to verify that SharePoint Server 2007 is supported when running inside Windows Server 2008 virtual machines on VMware vSphere 4.0. You start by selecting Office SharePoint Server 2007 from the product's drop-down list, as shown in Figure 2.49.

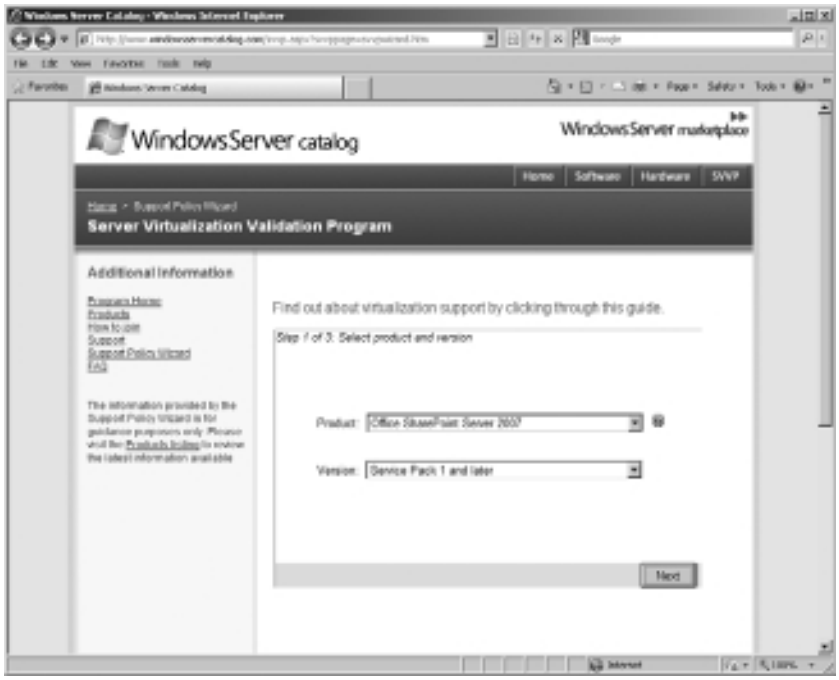

FIGURE 2.49 SVVP Support Policy Wizard, selecting the product

Next, you select VMware vSphere 4 as the virtualization platform and 64-bit Windows Server 2008 as the guest operating system, as shown in Figure 2.50.

FIGURE 2.50
SVVP Support Policy Wizard, selecting the virtualization platform and guest OS

The Support Policy Wizard shows that Office SharePoint Server 2007 running inside 64-bit Windows Server 2008 virtual machines on VMware vSphere is supported by Microsoft, as shown in Figure 2.51.

FIGURE 2.51
SVVP Support Policy Wizard support statement

For non-SVVP-validated configurations of ESX and Microsoft applications, support options vary depending on how customers purchase VMware and Microsoft products. For customers with a valid Microsoft Premier-level support agreement, Microsoft states that it will use "commercially reasonable efforts" to support its products running on VMware virtual machines. Server resellers including Dell, Fujitsu, Fujitsu-Siemens, HP, IBM, and Unisys offer end-to-end support for Microsoft software running on their servers and VMware if VMware products are purchased with the server hardware and are covered by a valid support agreement with the server reseller. This provides customers with one-stop support via the server reseller if an issue arises. For unusual or extremely difficult cases where direct engineering engagement is needed, VMware support customers can benefit from escalation through VMware's Premier support relationship with Microsoft.

In general, Microsoft offers its customers excellent support for its products running on VMware. Customers routinely tell us they receive the Microsoft support benefits that are effective and appropriate to maintain operations as planned.

Deployment Strategies

When you are ready to deploy a virtualized environment, you need to decide how the virtual machines in a solution stack should be placed in the vSphere cluster. This deployment discussion typically includes the following topics:

- Scale-up or scale-out strategy
- Virtual machine placement within the vSphere cluster

Scale-up or Scale-out Strategy

A *scale-out* approach uses multiple small virtual machines (for example, one vCPU per virtual machine), while a *scale-up* involves a large virtual machine with up to eight vCPUs per virtual machine, for example. The scale-out approach is generally more appropriate for virtual deployments using VMware vSphere.

The scale-out approach makes it easier to customize the virtual machine and application configuration as needed. Once you lock down the configuration, this approach usually provides better workload and security isolation and more granular change management. The scale-out approach also works very well for horizontal scalability, load balancing, and high availability, especially when applications use Windows Network Load Balancing or hardware load balancers.

VMware DRS functions more effectively with smaller virtual machines, with the added benefit of faster VMotion migrations. Smaller virtual machines can be easily migrated by DRS in order to balance the workload across the cluster. Virtual machines with more memory or more vCPUs put more constraints on DRS decision making.

A scale-up approach works well for applications such as SQL Server that can utilize the resources provided. Other than databases, most other server applications do not scale up very well and exhibit higher performance overheads associated with SMP scaling.

Virtual Machine Placement in vSphere Cluster

Another commonly asked question is, how should you place the various virtual machines in a solution stack within a vSphere cluster? For example, should you dedicate all SQL Server virtual machines to a single or a group of ESX servers, or should the web front-end server, SharePoint

Server, and SQL Server in a farm share a single ESX server? We recommend that you enable VMware DRS and let it optimally balance the virtual machines running on the ESX servers. The continuous real-time monitoring performed by DRS not only helps in the initial placement of virtual machines but also performs the necessary VMotion migrations to maintain balanced resource utilization across ESX clusters on an ongoing basis.

In general, do not consolidate multiple virtual machines that bottleneck on similar resources on the same vSphere host. This usually means spreading out workloads with similar profiles such as SQL Servers. Placing all web servers on a single ESX host can also lead to a single point of failure, negating any benefits of load balancing and high availability. To avoid all redundant virtual machines in a server farm from showing up on the same vSphere host at the same time, you should configure the DRS anti-affinity rules.

Mixing workloads that stress different types of resources usually balances out ESX server hardware resource utilization relatively better. If the components of the application stack communicating with each other are placed on the same ESX server, you also get the benefit of all network communication happening in memory and not having to go through physical network adapters. VMware studies have shown that this in-memory network traffic can exceed the typical 1 Gbps physical network bandwidth.

Summary

Windows-based servers account for the majority of x86 servers being deployed in data centers. As organizations increasingly adopt virtualization, they need guidance and best-practice recommendations for deploying their Windows server applications on the most preferred virtualization platform today, VMware vSphere. In this chapter, we discussed the topics that are common across these Windows server technologies and applications when virtualizing with VMware vSphere.

Hosted virtualization products such as VMware Workstation and Server are best suited for training, demo, or development environments. When you are ready to take virtualization to the next level and deploy enterprise applications in a virtual environment, using a bare-metal hypervisor is a must. Bare-metal hypervisors such as VMware ESX and ESXi provide direct access to the underlying hardware, thus allowing for greater performance and scalability.

Microsoft has updated its software licensing policies for server applications to make it more virtualization friendly. New Microsoft policies allow unlimited mobility of a large number of applications under defined conditions, making it easy to leverage key enhancements in virtual environments such as VMotion. This has opened the door to being able to fully utilize VMware DRS for automated workload balancing. These features along with VMware HA will provide the availability required by most organizations' tier-1 applications. If your organization has more stringent service-level agreements or requires the ability to use an application-aware high-availability solution, you can continue to do so in a vSphere environment. Microsoft Clustering Services and failover clusters are still possible and fully supported.

Deploying in a vSphere environment provides consolidation of specific operational procedures such as backing up virtual machines using VMware Data Recovery and patching virtual machines using VMware Update Manager. We discussed how you can leverage VMware Converter as well templates and cloning features to create virtual machines as well as for ongoing provisioning. Because virtual deployments are shared environments with multiple virtual machines consolidated on a single vSphere server, you must properly size the hosts, the network, and the storage

for your virtual deployment. You can leverage VMware's Capacity Planner, which can help with your initial virtualization assessment exercise.

Microsoft offers its customers excellent support for its products running on VMware. The Support Policy Wizard published as part of Microsoft's Server Virtualization Validation Program can provide support details for specific application versions you are interested in. A few of these applications, specifically those considered "tier-1" applications, are covered in the chapters to follow. As you will learn, deploying these applications on vSphere doesn't have to be intimidating or full of uncertainty. You will learn what to look out for and pitfalls to avoid, and you'll embark on your journey into virtualization with the confidence that you expect to have when deploying proven solutions.

Chapter 3

Virtualizing Windows Server 2008

With virtualization establishing roots in data centers across the world, the time will come when systems and application administrators will be asked to design and deploy tier-1 applications in a virtual environment. Such tier-1 applications include Microsoft SQL, SharePoint, and Exchange, among others. The common factor among these applications is the requirement to run on the Microsoft Windows Server platform.

This chapter will focus on running the Microsoft Windows Server 2008 operating system as a *guest*—the operating system running within the virtual machine—in a vSphere environment. Although the basic process of installing and configuring Microsoft operating systems in a virtual environment closely resembles the physical environment, you must take certain requirements into consideration. We will explore these requirements, take a look at the OS compatibility matrix, and review the features available to you when running in the vSphere environment.

With a basic understanding of the compatibility and technical requirements, you can begin working with virtual machines. Using examples, we will show how to build a virtual machine capable of running Microsoft Windows Server 2008 and how to install the operating system to establish a foundation for your applications. Once you've created and configured a working virtual machine to your required specifications, you can leverage one of the most timesaving features of running in a virtual environment: deployment via templates. By using templates, you can cut the time required to deploy new instances of Windows Server and ensure a consistent image is used throughout your virtual environment.

We will wrap up the chapter by covering a few methods of protecting your applications and data. Delivering high availability at the application level is not lost in the virtual environment, but you do need to ensure you conform to certain guidelines that have been tested and proven in the field. Microsoft failover clusters offer protection from hardware, OS, and application-level failures and are as capable of protecting your resources in vSphere as in a physical environment. We will review the requirements for building a supported Microsoft failover cluster in vSphere and show how to build a proof of concept. We will follow this up by covering data protection. A virtual machine is nothing more than a set of files on disk. This allows you to encapsulate all files that make up a virtual machine for easy backup, transport, and recovery. Additionally, you have the option of running traditional backup and recovery software to meet the needs of your organization.

At the end of this chapter, you will have gained useful knowledge that will ease your journey into the virtual world. Let's begin this journey with the basic requirements needed to get off the ground.

Understanding Guest OS Compatibility and Requirements

Before you start down the road of loading CDs (or rather DVDs these days) for your guest operating system (GOS or guest OS), you need to consider system compatibility and requirements. Software vendors generally do a good job of letting end users know what software editions and versions they have tested with their products. Additionally, you as a user have the responsibility to follow the guidelines and requirements set by the vendors to ensure smooth deployments of these products. This is important both for the vendor and for you as a customer. Vendors have to provide support for their products across many differing environments, and not having clearly defined support and compatibility matrices would result in an unsupportable mess of varying configurations.

Guest OS Compatibility

The guest OSs that are supported for running under vSphere have been through VMware's testing processes and are validated by VMware to run as guest OSs. This is an ongoing process as new operating system builds go into beta, release candidate, and eventually general availability. Naturally there is some lag between the time when a new version of an operating system is released to when it can become an officially supported guest operating system. However, it is not uncommon for "experimental" (or limited) support to be available in the interim. This level of support allows for limited testing of the new guest OS within your environment as long as you are aware that some features such as guest OS deployment and customization may not work. Additionally, drivers for VMware virtual hardware such as VMXNET Ethernet adapters and VMware paravirtualized SCSI controllers may not be compatible while the guest OS is under "experimental" support.

As of vSphere 4 Update 1, the following Windows 2008 editions are fully supported by VMware to run as a guest OS on ESX and ESXi:

- Windows Server 2008 Web Server (SP1, SP2: 32/64-bit; R2: 64-bit)
- Windows Server 2008 Standard edition (SP1, SP2: 32/64-bit; R2: 64-bit)
- Windows Server 2008 Enterprise edition (SP1, SP2: 32/64-bit; R2: 64-bit)
- Windows Server 2008 Datacenter edition (SP1, SP2: 32/64-bit; R2: 64-bit)
- Windows Server 2008 Small Business Server Standard/Premium (SP1, SP2: 64-bit)
- Windows Server 2008 Essential Business Server Standard/Premium (SP1, SP2: 64-bit)

For the latest updates to the support matrix, visit the VMware website: `www.vmware.com/resources/compatibility`.

Guest OS Requirements

In general, the requirements set forth by the OS vendor will remain the same when virtualizing the operating system. This holds true for Windows 2008 and Windows 2008 R2. Table 3.1 and Table 3.2 list the minimum requirements for Windows 2008 and Windows 2008 R2, respectively, for installing these as guest operating systems in vSphere 4 and vSphere 4 U1.

TABLE 3.1: Windows 2008 Requirements

COMPONENT	REQUIREMENTS
Processor	Minimum: 1 GHz (x86) or 1.4 GHz (x64)
	Recommended: 2 GHz or faster
Memory	Minimum: 512 MB RAM
	Recommended: 2 GB RAM
	Maximum (32-bit): 4 GB (Standard), 64 GB (Enterprise and Datacenter)
	Maximum (64-bit): 8 GB (Foundation), 32 GB (Standard), 2 TB (Enterprise and Datacenter)
Disk space	Minimum: 20 GB (32-bit), 32 GB (64-bit)
	Recommended: 40 GB

Note: Virtual machines with greater than 16 GB of RAM will require more disk space for paging, hibernation, and dump files.

TABLE 3.2: Windows 2008 R2 Requirements

COMPONENT	REQUIREMENTS
Processor	Minimum: 1.4 GHz (x64)
	Recommended: 2 GHz or faster
Memory	Minimum: 512 MB RAM
	Recommended: 2 GB RAM
	Maximum: 8 GB (Foundation), 32 GB (Standard), 2 TB (Enterprise and Datacenter)
Disk space	Minimum: 32 GB (64-bit), 10 GB for Foundation
	Recommended: 40 GB

Note: Virtual machines with greater than 16 GB of RAM will require more disk space for paging, hibernation, and dump files.

Deploying Windows 2008 Virtual Machines

With an understanding of the compatibility considerations and system requirements you must adhere to, you can begin the process of deploying Windows 2008 as a guest operating system in your vSphere environment. Installing the operating system from media is not the first step, but it's not far from it. The first objective is to establish some working ground for your OS. This will come in the form of a virtual machine.

Your virtual machine will consist of many components you are familiar with such as CPU, memory, disks, CD-ROM, and more. It's your job to build a virtual machine using the recommended configuration based on your guest OS requirements and recommendations from VMware. We will cover the various types of virtual hardware available to you and the recommended configurations. Once you have an idea of what the virtual machine will consist of, you can commit these changes to the VM and begin the process of installing the guest OS. The desired result is that you have a working base install of a virtual machine and guest OS that you can then use as a template.

A *golden template* is key to optimizing your deployment strategy and having a consistent baseline to build upon. We will cover the most basic approach for building a template that you will be able to use for most application deployment scenarios. With your template ready to go, you will create a guest customization specification; think of this as Microsoft Sysprep integrated into vCenter Server, for a one-stop deployment. Guest customizations such as Sysprep allow you to set the attributes in the guest OS that make that system unique on the network.

Although most of your deployments in the virtual environment will be in the form of new virtual machines, you do have the option to clone a physical system and convert it into a virtual machine. This is especially helpful for systems that have been inherited that you may not be truly familiar with, for systems that require extensive application configuration or tuning, or for quickly reducing your physical footprint. We will cover the various components of the VMware vCenter Converter tool to simplify the process and outline the process of converting a physical system.

Choosing Virtual Hardware

This chapter really does you no good without a foundation on which to build your guest operating system. For us, this foundation is the virtual machine and its hardware. vSphere 4 has added some great new features, improved virtual hardware performance, and increased the capacity of much of this hardware. As a result, you are now able to scale your Windows virtual machines higher than ever before. Windows 2008 is able to take advantage of all of these improvements, especially those such as hot-add CPU and memory. This is huge! Virtualization made it easy to deploy a system as "thinly" as possible in the past because you could always take the machine down and add resources if you needed. Now you are able to deploy those "thin" systems and not incur any downtime when you need to add memory or virtual CPUs. Of course, this may be information you want to keep close to the vest—you know how excitable your customers will be once they know that you have the ability to do this.

We'll cover the virtual hardware available to you for all operating systems, and we'll break down what you, as a Windows 2008 administrator, will want to focus on.

As you can see in Table 3.3, you have quite an extensible platform on which to build up the virtual machine to suit your needs. Note the maximum configurable quantities for each type of virtual hardware available.

VMware has done a great job at making the default selections compatible with most guest operating systems. This allows you to build a new VM and install the guest OS without requiring extra driver disks. In most cases, this gets you to a running state after which you can install VMware-specific drivers to obtain better performance.

Table 3.4 shows the various types of virtual hardware available for your virtual machines and specifically whether each type is supported and/or recommended for Windows 2008.

TABLE 3.3: Virtual Machine Hardware

COMPONENT	QUANTITY
CPU	1 (2-8 with VMware SMP)
Memory	255 GB
Video controller	1
Floppy controller	1
Floppy drives	2
IDE controller	1
CD-ROMs	4
SCSI adapters	4
SCSI devices	1–15 per adapter

TABLE 3.4: Virtual Hardware Support

NETWORK	SUPPORT	DETAILS
e1000	Supported	Emulated Intel 8254EM 1Gbps NIC
Enhanced VMXNET	Supported	Second-generation VMware Virtual NIC
Vlance	Not Supported	Emulated AMD 79C970 PCnet32 Lance NIC
VMXNET	Not Supported	VMware Virtual NIC
VMXNET 3	Supported	Third-generation VMware Virtual NIC
STORAGE	**SUPPORT**	**DETAILS**
BusLogic	Not Supported	Virtual BusLogic Parallel SCSI adapter
IDE	Supported	Virtual IDE Adapter for ATA disks
LSI Logic	Supported	Virtual LSI Logic Parallel SCSI adapter
LSI Logic SAS	Supported	Virtual LSI Logic SAS adapter
VMware vSCSI	Supported	VMware virtual SCSI adapter

TABLE 3.4: Virtual Hardware Support *(CONTINUED)*

VIRTUAL HARDWARE	SUPPORT	DETAILS
Hot-add memory	Supported	Add memory to running virtual machine
Hot-add vCPU	Supported	Add a virtual CPU to running virtual machine
Symmetric multiprocessing	Up to eight vCPUs	Maximum number of virtual CPUs

As you can see from Table 3.4, you have a few choices to make when creating your virtual machine. This can be a bit overwhelming if it's your first time deploying a virtual machine. Sure, the virtual hardware is supported and you may even have the drivers built in, but is it the right selection for what you're trying to accomplish or a selection that VMware even recommends? Luckily for you, the virtual machine creation process takes care of most of this for you, as you'll see in the next section.

Creating the Virtual Machine

Well, now you know what you can and can't use as far as virtual hardware is concerned, regardless of whether it's the right thing to do … we'll get to that soon enough. By planning the first few builds correctly, you will be able to avoid spending hours creating new virtual machines and instead use the portable nature of virtual machines to your advantage. Let's get down to the nitty-gritty and start building some virtual machines!

In this first demo, you will build a virtual machine capable of running Microsoft Windows 2008 R2 Enterprise edition. In most cases, the actual virtual machine will be the same between different versions of the OS (that is, Standard vs. Enterprise), but if there are any differences, we'll mention them along the way. We will be using an ISO image of the install media, but if you have physical access to the ESX host where the virtual machine will be built, you can also use physical media to install the OS by placing the install media in the optical drive on the ESX host.

NOTE An ISO image is a disc image of an optical disc in the format defined in the ISO 9660 file system standard. This is the same file system used with CD-ROM media. These images are a great way to make backups of your physical media, they are portable, and luckily for you they are mountable by your virtual machines. This gives you the ability to copy your install media (ISO image) to the ESX host, configure the virtual machine to mount the CD-ROM device at start-up, and install the guest operating system without having to fumble around for CDs or DVDs.

Most CD/DVD-burning programs allow you to burn to a file, and from there you can choose the type of file (`*.iso`). If you do not have burning software handy, a quick web search will yield a few free results that will get you going in the right direction.

1. Launch the VMware vSphere client from the Start ➢ Programs menu, enter the hostname or IP address of your vCenter server (or ESX host), and enter your username and password to log on (Figure 3.1).

FIGURE 3.1
vSphere client login screen

2. You have a few options as to where you initiate the creation of the virtual machine. For the sake of this demo, you will create the VM at the host. Right-click the host, and select New Virtual Machine.

3. The Create New Virtual Machine Wizard gives you the option to use a typical or custom configuration (Figure 3.2). Select Custom, and click Next.

FIGURE 3.2
Getting started configuring your new virtual machine

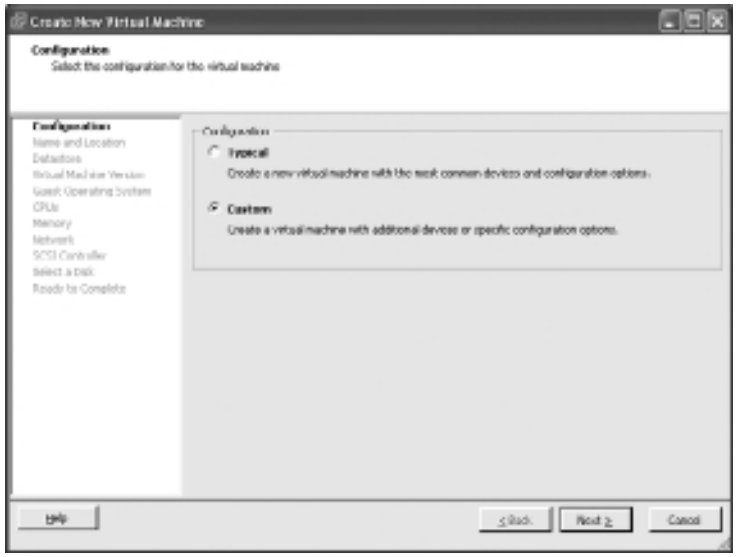

4. On the Name And Location page, enter a name for your virtual machine (Figure 3.3). It is typically a good idea to name the VM the same as the guest OS hostname. Doing so will save you time (and grief) when searching for or trying to identify VMs. The name given here is also used to name the folder and files that make up the VM.

   Under Inventory Location, you can select the logical placement of the virtual machine.

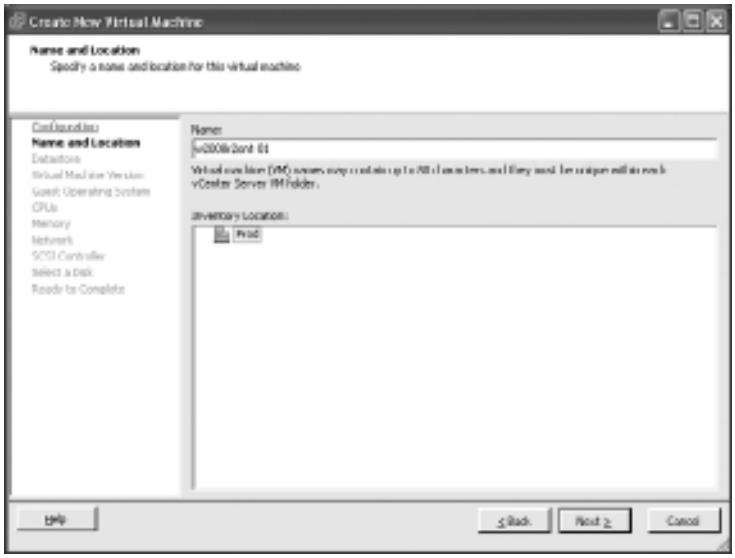

FIGURE 3.3
Choosing a name and location for your VM

5. The Datastore selection page allows you to select the appropriate datastore for the system/boot disk (Figure 3.4). If your ESX host has local storage, you will also see this volume listed. If you plan to use any of the advanced features of vSphere (VMotion, HA, or DRS), you will want to place your VMs on shared storage. Select the datastore to place the virtual machine files, and click Next.

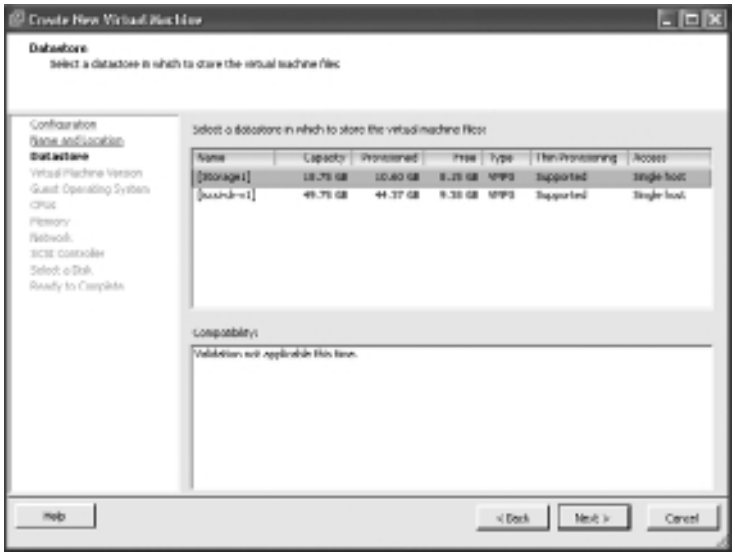

FIGURE 3.4
Selecting an appropriate datastore for your virtual machine

6. When creating a virtual machine in vSphere, you have the option to create either a version 4 or version 7 virtual machine. Version 4 virtual machines are compatible with ESX version 3.0 and newer and VMware Server 1.0 and newer. Version 7 virtual machines are compatible with ESX version 4.0 and newer and VMware Server 2.0. Select Virtual Machine Version: 7, and click Next (Figure 3.5).

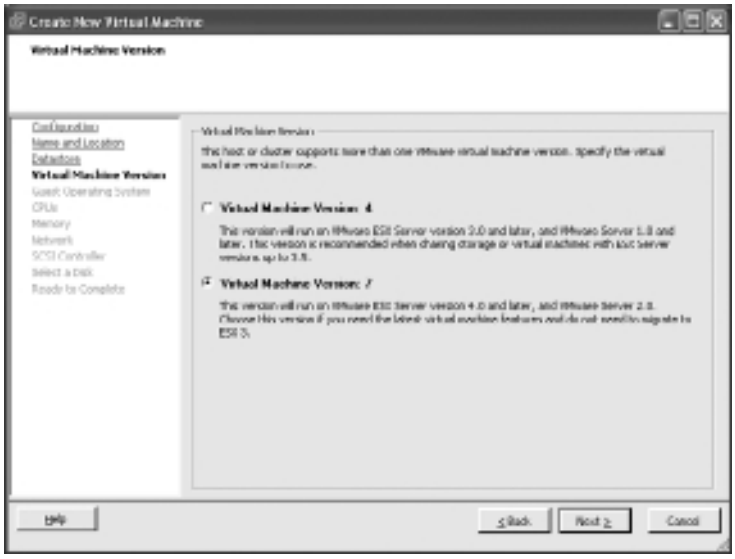

FIGURE 3.5 Specifying the virtual machine version based on the ESX host version

NOTE Virtual hardware version 7 was introduced with vSphere and is available to virtual machines created or migrated to ESX 4.0 hosts and higher. Virtual hardware version 7 introduces new SCSI controllers (LSI Logic SAS and VMware paravirtualized SCSI), a new virtual Ethernet controller (VMXNET 3), and a new virtual graphics card.

7. The Guest Operating System page lists the supported operating systems and versions for which you can create the virtual machine (Figure 3.6). The guest OS and version you select will trigger the default virtual hardware to be automatically selected. This includes the amount of memory, the type of virtual NIC, and the SCSI adapter.

   Select Microsoft Windows Server 2008 R2 (64-bit), and click Next.

WARNING In vSphere versions prior to 4.0 Update 1, Windows Server 2008 R2 and Windows 7 will be listed as experimental. These guest OS versions should be deployed on vSphere 4.0 Update 1 or newer to be fully supported.

WARNING If you select a version of a guest operating system and choose to change it before completing the wizard, close the wizard and start over. In the 4.0 and 4.0 Update 1 client, simply changing the version of the GOS does not update the default virtual hardware for the GOS version being deployed.

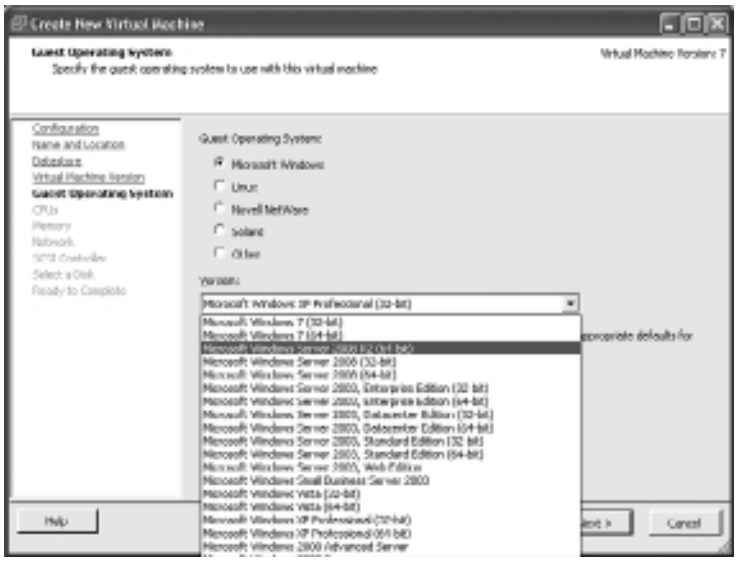

FIGURE 3.6
The Guest Operating System page helps determine virtual hardware settings and recommendations.

8. In previous versions of Windows, deploying a single vCPU virtual machine installed the uniprocessor HAL by default. Increasing the number of virtual CPUs at a later time would require the installation of the multiprocessor HAL. With Windows 2008, the multiprocessor HAL is installed by default, allowing for the addition of CPUs without the need to update drivers.

   Select the number of CPUs (one to eight depending on licensing and guest OS support) based on the requirements for the virtual machine, and click Next (Figure 3.7).

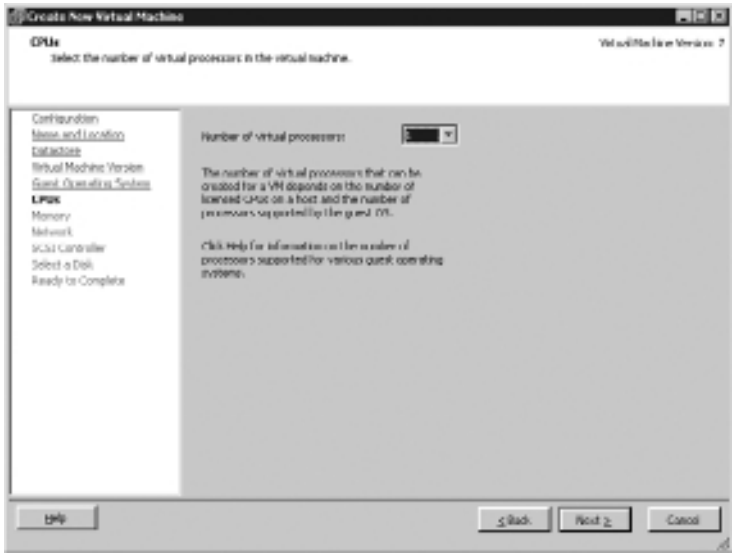
Create New Virtual Machine
CPUs
Select the number of virtual processors in the virtual machine.
Configuration
Name and Location
Datastore
Virtual Machine Version
Guest Operating System
CPUs
Memory
Network
SCSI Controller
Select a Disk
Ready to Complete
Number of virtual processors:
The number of virtual processors that can be created for a VM depends on the number of licensed CPUs on a host and the number of processors supported by the guest OS.
Click Help for information on the number of processors supported for various guest operating systems.
Help
< Back
Next >
Cancel

FIGURE 3.7
Selecting the number of CPUs

9. The Memory allocation page will be set to the default recommended value for the guest OS selected (Figure 3.8). You will also be given the minimum, optimal, and maximum recommendations for your guest OS. Leave the default recommended value (4 GB for Windows 2008 R2), and click Next.

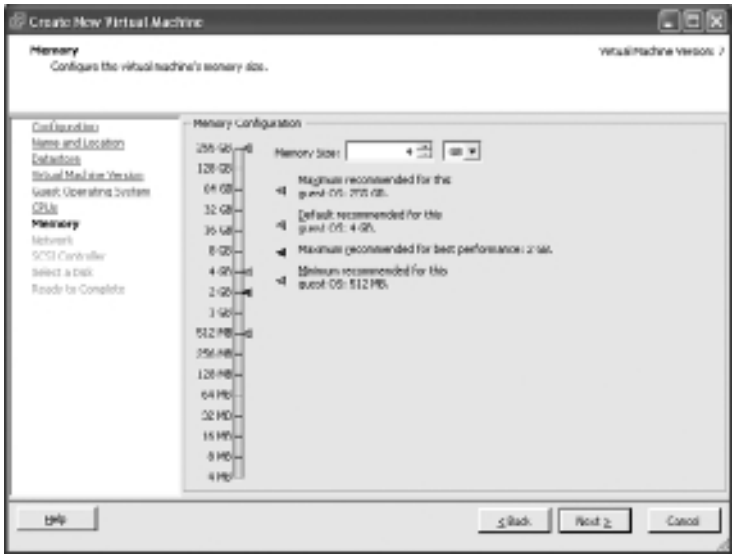

FIGURE 3.8
Memory allocation with recommendations based on guest OS chosen

10. On the Network connection page, you have the opportunity to add up to 10 virtual Ethernet devices (Figure 3.9). The default adapter type for Windows 2008 R2 is the e1000 Ethernet adapter. Selecting this adapter will ensure you are able to gain network access without needing to install additional drivers. This is handy if you want to perform a domain join or perform other network activities before you have a chance to interact with the guest.

    If you have the option to wait until the guest OS is installed, you can select the recommended Ethernet adapter, VMware VMXNET3. Once the VMware Tools package is installed, you will be able to use the adapter. For the sake of this demo, you will install one adapter, the default adapter (e1000); select your desired network (VM Network); and make sure the Connect At Power On box is selected.

11. The SCSI Controller page lists the available virtual SCSI controllers, and as before, the default recommended choice is preselected. For Windows 2008 and Windows 2008 R2, you have the option to select the LSI Logic parallel, LSI Logic SAS (default recommended), or VMware paravirtual SCSI controllers.

TIP vSphere 4 Update 1 introduced the ability to use the VMware paravirtual SCSI (PVSCSI) controller as the boot controller. Prior to this, using the PVSCSI controller as the boot device was not supported.

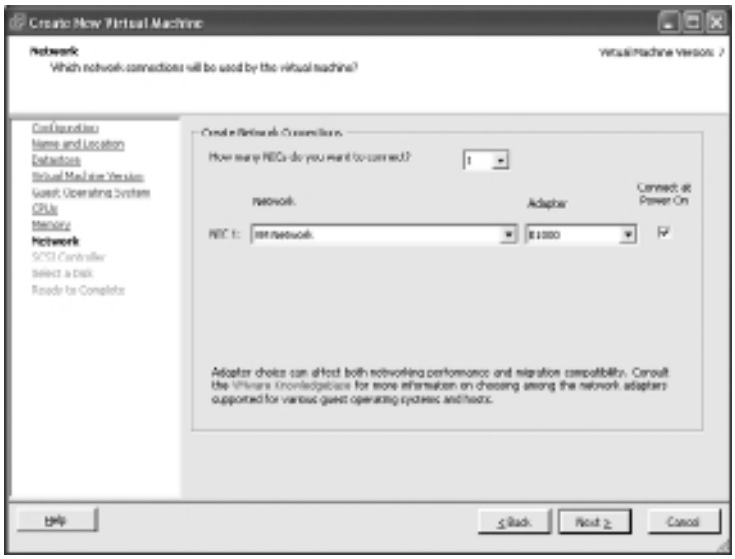

FIGURE 3.9
Selecting the number of NICs, virtual network port group, adapter type, and post power-on connectivity

Both LSI Logic drivers are built into Windows 2008 and Windows 2008 R2, thus allowing you to install the guest OS without requiring a driver disk.

Leave the default LSI Logic SAS controller selected, and click Next (Figure 3.10).

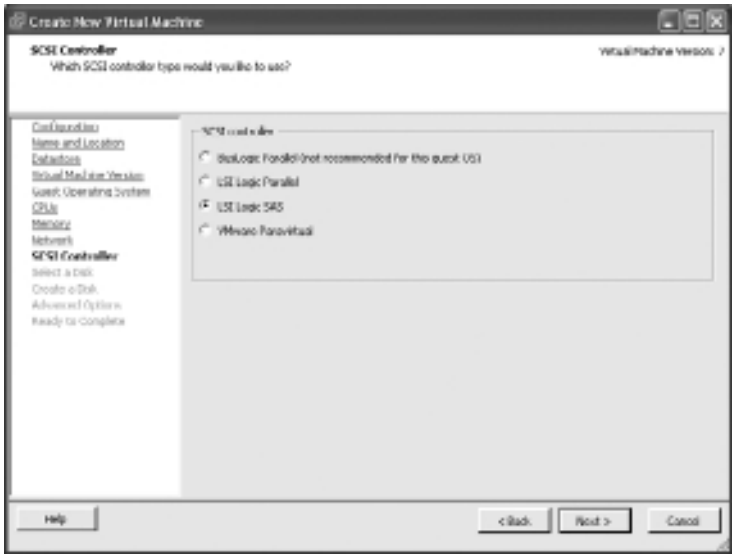

FIGURE 3.10
Select the SCSI controller; the default selection is the recommended choice.

12. Next you must select whether you will be creating a new virtual disk using an existing virtual disk, using a raw device mapping, or using no disk at all (Figure 3.11). Because this is a new install, you will choose to create a new virtual disk and continue to the next screen.

Create New Virtual Machine
Select a Disk
Virtual Machine Version: 7
Configuration
Name and Location
Datastore
Virtual Machine Version
Guest Operating System
CPUs
Memory
Network
SCSI Controller
Select a Disk
Create a Disk
Advanced Options
Ready to Complete
A virtual disk is composed of one or more files on the host file system. Together these files appear as a single hard disk to the guest operating system.
Select the type of disk to use.
Disk
Create a new virtual disk
Use an existing virtual disk
Reuse a previously configured virtual disk.
Raw Device Mapping
Give your virtual machine direct access to SAN. This option allows you to use existing SAN commands to manage the storage and continue to access it using a datastore.
Do not create disk
Help
< Back
Next >
Cancel

FIGURE 3.11
Choosing whether to create a new virtual disk, use an existing virtual disk, or not create a disk

Before the actual virtual machine disk file (VMDK file) is created, you must specify the size, location, and provisioning mechanism for your virtual disk.

The capacity will be the maximum amount of space allocated to the virtual disk file and what your guest OS will see as the capacity of the boot disk. The default (40 GB for Windows 2008) is again prepopulated depending on the guest OS you chose in step 7.

Disk provisioning determines how the disk will be created. If left blank, all of the disk space you specified in the capacity section will be provisioned and committed at once against physical storage. Selecting the Allocate And Commit Space On Demand (Thin Provisioning) box will allow for the overcommitting of physical disk space by growing the virtual machine disk file as new blocks are written, up to the maximum size specified. This is the default provisioning method when placing the virtual machine disk file on NFS storage.

If thin provisioning is selected, the Support Clustering Features Such As Fault Tolerance box is grayed out, and vice versa. Fault Tolerance (and other clustering technologies such as Microsoft failover clusters) will require that the virtual machine disk file is created as eagerzeroedthick. Eagerzeroedthick will fully provision and commit all space to the VMDK file as well as write zeros to all blocks of the disk file. This improves performance on the first write at the expense of taking longer to provision the disk during deployment.

WARNING **When choosing to thin provision production systems, it is important to have proper monitoring and alerting set up using either vCenter Server alarms or a third-party tool for monitoring storage.**

The default location for the virtual machine disk file is with the other virtual machine files. This is usually the best choice because it keeps the boot disk and the core virtual machine files (`.vmx`, `.vmsd`, `.nvram`) together. However, it can be placed on any available datastore:

1. Keep the defaults for the size and disk file location, but because better performance is usually not a bad thing, select the Support Clustering Features Such As Fault Tolerance box (Figure 3.12), and click Next.

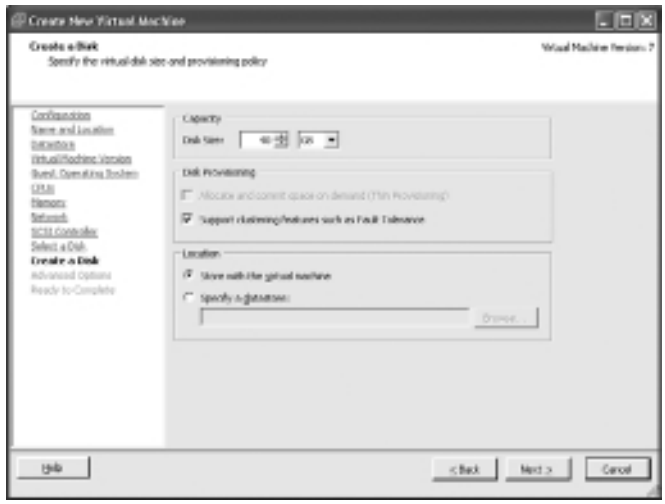

FIGURE 3.12
Selecting the desired capacity of the new virtual disk, whether to enable thin provisioning and/or FT, and where to store the VMDK file

2. The next page will usually be left with the default settings for a new build with a single boot disk (Figure 3.13). The Advanced Options page lets you choose the SCSI (or IDE) bus and target number and whether the disks should be independent and persistent or nonpersistent. We'll dig a little deeper into what these are and when you can use them a bit later in the chapter.

   For this example, accept the default setting, which is SCSI (0:0), and click Next.

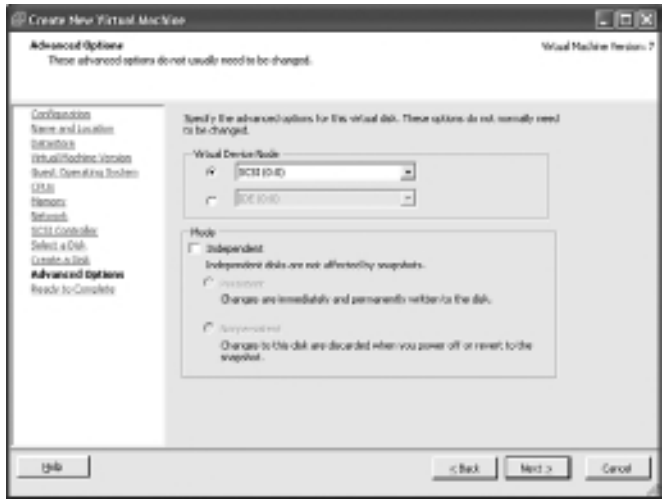

FIGURE 3.13
Using Advanced Options to select a specific SCSI bus and target setting and whether the virtual disk is independent and persistent or nonpersistent

3. The Ready To Complete page summarizes the settings you have configured and gives you an opportunity to edit the virtual machine before the first power on. This will save you a few clicks because you will have to attach an ISO image of the Windows 2008 R2 install media before powering on. This also gives you an opportunity to check some of the other settings before continuing.

Select the Edit The Virtual Machine Settings Before Completion box, and click Continue (Figure 3.14). Clicking Continue will present the proposed virtual machine hardware and give you a chance to make changes not available through the wizard (see Figure 3.15).

FIGURE 3.14
Ready To Complete summarizes the virtual hardware and allows for editing additional settings before creation.

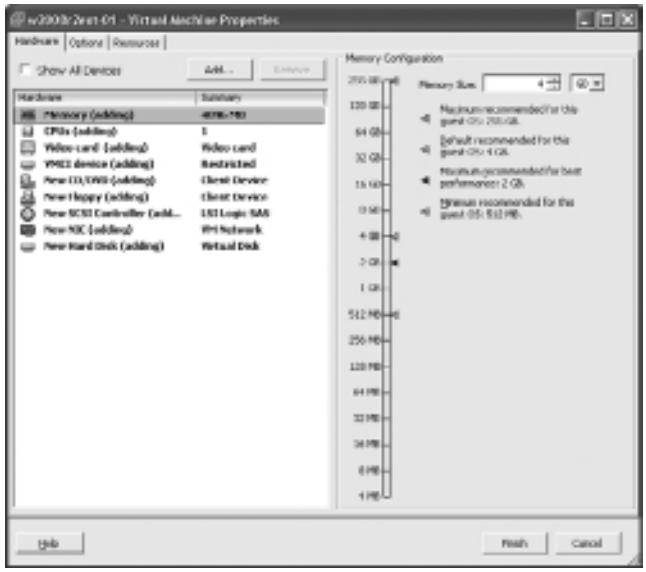

FIGURE 3.15
Editing the virtual machine settings prior to creation allows for connecting devices such as install media or disconnecting devices such as network interfaces.

4. With the virtual machine properties open, you can select the CD/DVD virtual device and mount the ISO image of your install media (Figure 3.16). Select the Connect At Power On box under Device Status. Click Finish to complete the VM creation.

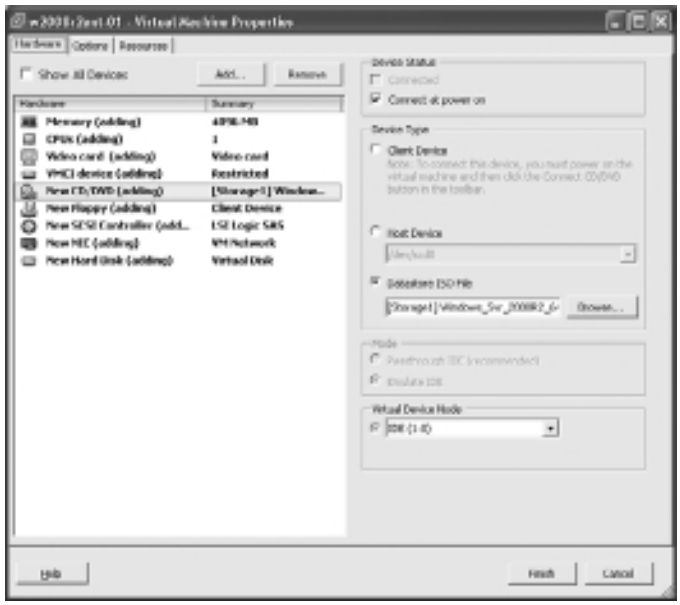

FIGURE 3.16
Attaching the ISO image prior to VM creation will save a few clicks prior to powering on the VM.

TIP If you have access to the physical ESX host and physical install media, you could select the Host Device radio button and load the install media to install directly from the disc.

Once you have created your ISO image, you will have to copy it to a location accessible from your ESX host. You can copy the ISO image directly to the ESX host and place the file either on a VMFS datastore or in the `/vmimages` directory. To copy the file directly to the VMFS datastore, you can use the datastore browser. To copy the file to the `/vmimages` directory, you may use a SCP tool (Secure Copy) such as `winscp.exe` (`http://winscp.net/`).

5. Locate and select the VM from the inventory. The Summary tab (Figure 3.17) gives an overview of the resources allocated, network and storage connectivity, and resource utilization when the VM is running.

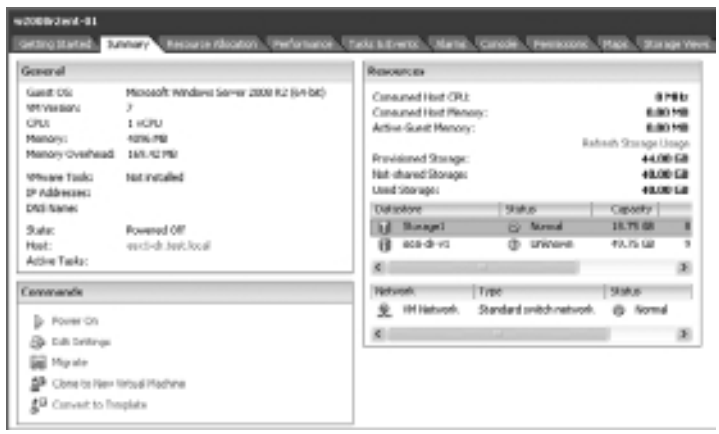

FIGURE 3.17
Review the Summary tab for a quick look at the VM configuration, current host, and resource usage.

Installing Windows 2008 R2

The virtual machine you've now built is the basic framework for what could be tens, hundreds, or even thousands of virtual machines you may build into your virtual environment. With this framework laid out, you can now get your operating system installed. How you go about this is mostly a matter of choice, but in most cases you will either mount an ISO image to the virtual CD-ROM device (as previously mentioned) or physically load the media into the ESX host's optical drive.

If you are lucky and you do a good job of getting this OS installed and configured, you won't have to become a pro at installing Windows. Remember, the point here is to build a golden template that you will use as the base for all other new virtual machine deployments. Follow this procedure to install the OS, followed by some options for creating your deployable template:

1. With the VM selected, click Open Console (Figure 3.18).
2. Power on the VM by clicking the green arrow or selecting Power ➢ Power On in the VM menu. If everything has gone to plan, you will see the "Starting Windows" splash screen and finally the Windows installation dialog boxes.
3. Make any changes to the language and keyboard preferences, and click Next.
4. Click the Install Now button to kick off the interactive portion of setup.

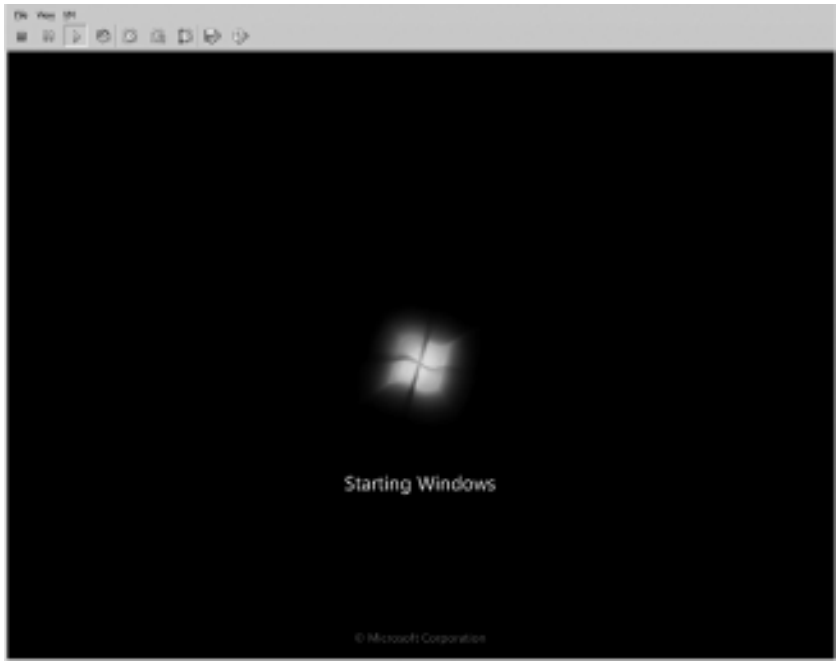

FIGURE 3.18
The console view gives a KVM-like display without the requirement for the OS to be running or accessible on the network.

Manually Choosing the Boot Device

Although not default behavior for a new virtual machine, you may be required to manually pick the boot device. If this is the case, the easiest method for doing so is to edit the virtual machine properties to add a power-on boot delay. This will give you time to hit Esc and select a boot device.

1. With the VM powered down, right-click the VM, and select Edit Settings.
2. In the virtual machine properties, under the Options tab, select Boot Options under Advanced, and enter a delay time in milliseconds. Because of the delay from the time the VM is powered on to when the console is actually displayed, we like to set this to a high number such as 15,000 or 30,000 milliseconds.
3. Click OK in the virtual machine properties. You can now open the console to the VM and power on.
4. When the BIOS screen appears, hit the Esc key, which will present the boot menu, as shown in the following illustrations.
5. You can then use the up and down arrows to select the boot device as CD-ROM Drive and continue with the setup.

Don't forget to set the boot delay back to 0 after setup is complete.

5. Depending on your media type, you may be presented with some options as to the version of Windows you would like to install; for this installation, we will choose the full installation of Windows 2008 R2 Enterprise and click Next.

6. After reading and agreeing to the licensing terms, select the box, and click Next.
7. The type of installation refers to whether you are upgrading an older operating system such as Windows 2003 or installing a fresh OS. Since you will be installing a new copy of Windows here, you will select Custom (Advanced).

The Where Do You Want To Install Windows? window (Figure 3.19) can be either welcoming or alarming. In most cases, if you have created the virtual machine using the defaults presented by vSphere, you will not have any trouble, and you will be able to simply select your disk, as shown in Figure 3.19, as Disk 0 Unallocated Space, and either click Next or choose to carve it up further. The recommendation is to keep the allocated space for the OS and deploy more data drives later if necessary.

FIGURE 3.19
Selecting a suitable location to install Windows

However, if you have chosen to select a SCSI controller for which Windows does not have a driver available, you will be presented with the blank screen shown in Figure 3.20.

Luckily, VMware has made this much simpler than in years prior. It used to be that you would have to launch a browser from your workstation, embark on a journey to find the correct drivers, and then figure out how you could get them into the VM during setup so that Windows could recognize the controller. Those days are behind you now.

FIGURE 3.20
When the Windows installer detects an unknown SCSI Controller, click the Load Driver link to select the driver location.

VMware now ships these drivers on floppy images available on the file system of ESX classic. Located under the `/vmimages/floppies` directory (Figure 3.21), you can select the file `pvscsi-1.0.0.5-signed-Windows2008.flp` file and attach it to the floppy drive of your virtual machine. Once connected, you can click Load Driver, which will prompt you to insert the driver disk (which you have done, virtually). Then click OK, which will prompt Setup to read the available drives and present you with the list of available drivers (Figure 3.22). You can then select the driver that fits your installation and proceed by clicking Next.

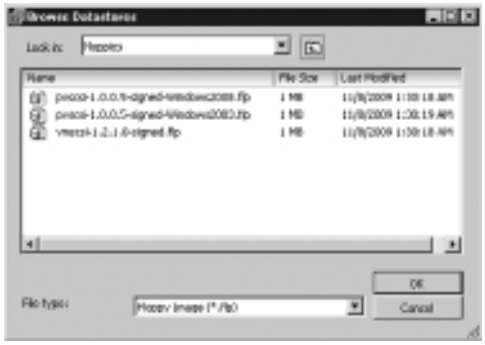

FIGURE 3.21
Selecting the appropriate driver disk image from the `/vmimages/floppies` directory on the ESX host

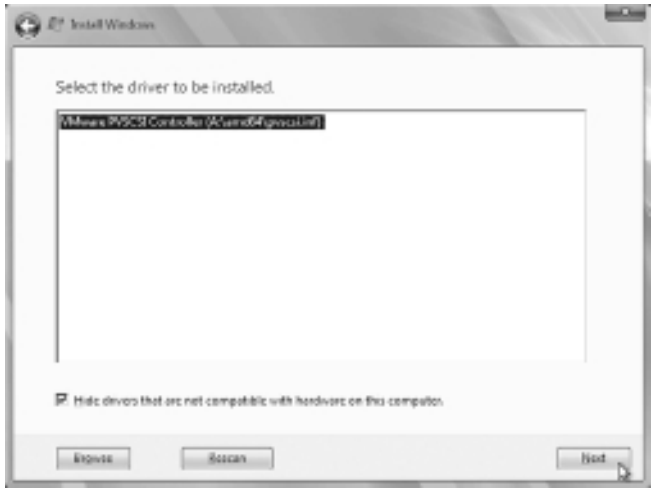
Select the driver to be installed.
Hide drivers that are not compatible with hardware on this computer.
Browse...
Rescan...
Next

FIGURE 3.22
The Windows installer will detect the new media and present the driver options.

Now would be the time to stretch or get a refill of coffee as you wait for the Windows 2008 R2 installation to complete. Once Setup has completed, you will be asked to enter a new password for the local administrator account.

Once logged in, you can kick off the installation of the VMware Tools package:

1. Right-click the VM, select Guest, and click Install/Upgrade VMware Tools.
2. Select Run Setup.exe when prompted; this will launch the VMware Tools installer (Figure 3.23).
3. Select a typical installation, click Next, and click Install.

AutoPlay
CD Drive (D:) VMware Tools
Always do this for software and games:
Install or run program from your media
Run setup.exe
Published by VMware, Inc.
General options
Open folder to view files
using Windows Explorer
View more AutoPlay options in Control Panel

FIGURE 3.23
The VMware Tools installer will auto-run when the option to install/upgrade VMware Tools is selected.

SERVER CORE INSTALLATIONS

If you choose to install the Server Core edition of Windows 2008 or Windows 2008 R2, you will not be prompted to run VMware Tools `Setup.exe`. Server Core editions of Windows 2008 do not include the `Explorer.exe` shell and must be managed by command line or remotely. To install VMware Tools in Server Core, right-click the VM, select Guest, and click Install/Upgrade VMware Tools. This will mount the VMware Tools ISO. At the command line, change the directory to D: and run `Setup.exe` by typing the following:

```
D:
Setup.exe
```

Follow the wizard to select a typical installation, and click Next and then Install.

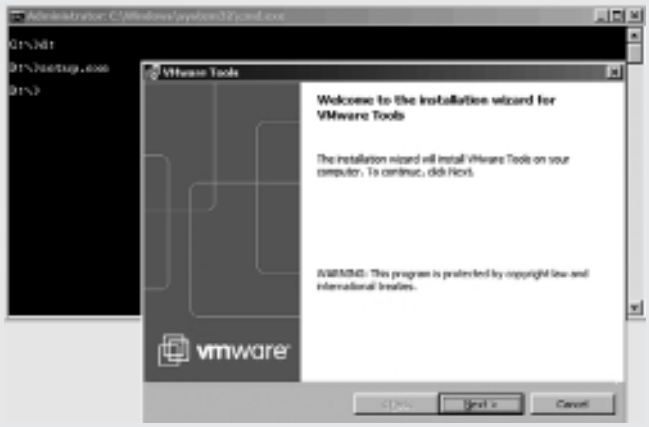
VMware Tools
Welcome to the installation wizard for VMware Tools
The installation wizard will install VMware Tools on your computer. To continue, click Next.
vmware

WARNING At the time of this writing, there was a VMware display driver issue with Windows 2008 R2. The VMware Tools installation package in vSphere 4 Update 1 does not install the VMware SVGA driver. If you are installing in versions prior to vSphere 4 Update 1, you can avoid this issue during installation by selecting the Custom installation type. When given the option, select the SVGA driver, and disable the installation of the drivers.

Once that's completed, verify hardware acceleration is set to full for the best performance when working through the console:

1. In the Control Panel, select Adjust Screen Resolution under Appearance.
2. On the Screen Resolution page, click Advanced Settings; this will open the adapter properties for the display adapter.
3. On the Troubleshoot tab, click the Change Settings button, and slide the Hardware Acceleration marker all the way to the right (Figure 3.24).

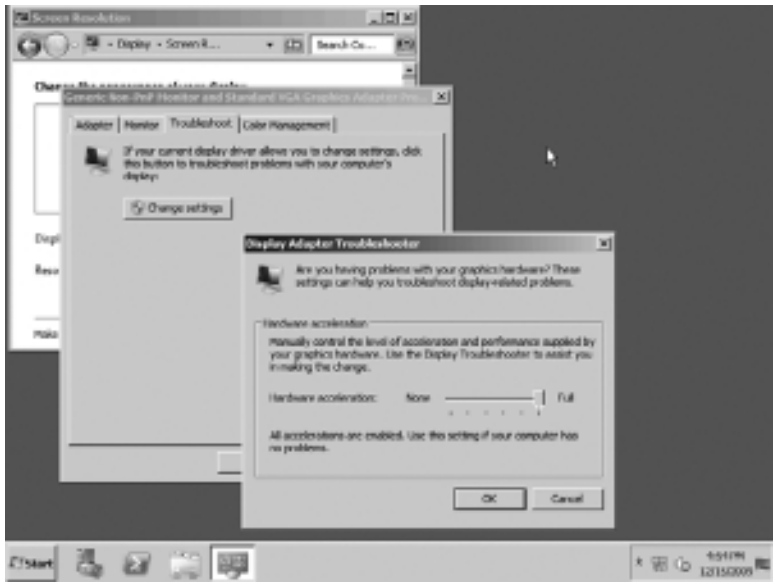

FIGURE 3.24 Hardware acceleration improves performance when working on the console through the vCenter client.

A reboot is required for the changes to take effect.

You now have a working installation of Windows, using the recommended virtual hardware and running VMware Tools. Many of your own internal procedures will still need to be followed before you could dare put this server into production. The following are the most basic procedures to perform before moving forward:

- Setting the time zone and time synchronization (see "Synchronizing Time")
- Installing antivirus software
- Installing the latest updates
- Enabling Remote Desktop
- Changing the screen saver to blank (no need to waste cycles on a screen saver in a VM)
- Disabling hibernation (`powercfg.exe -h off`)
- Setting the power plan to high performance
- Setting the system settings to adjust for best performance

Synchronizing Time

Timekeeping in virtual machines can be a bit tricky. You have to be aware of time drift that occurs in the guest operating system as a result of idle CPU cycles. As your virtual machine gets less CPU time, which is usually the case on less-demanding VMs, your clock drifts more. Because of this, you want to make sure you pay special attention to time synchronization within the guest. For Windows VMs, there are two recommended ways of keeping time under control—using the built-in w32 time service or using a third-party utility to query NTP servers.

If your guest OS will run as a domain controller, you must take additional steps to ensure that the time is accurate. For more information on configuring time synchronization for domain controllers, see Chapter 4, "Virtualizing Active Directory."

For guests that will be added as members of a domain but will not be domain controllers themselves, the recommendation is simple—disable VMware Tools time synchronization between the virtual machine and the ESX server. The domain time synchronization hierarchy (the default setting for Windows systems) should be followed so that time is served from the domain controllers. Synchronizing with any other time source can lead to time differences between the guest and the domain controllers that can in turn lead to login issues. Follow this procedure to disable VMware Tools time synchronization:

1. Open a console or Remote Desktop session to the virtual machine in question.
2. Double-click the VMware Tools icon in the system tray.
3. In the VMware Tools Properties dialog box, on the Options tab, verify the Time Synchronization Between The Virtual Machine And The ESX Server box is deselected.

If the guest OS will be a stand-alone server, in other words, not added to a domain, you can follow the next procedure to make the recommended changes for proper timekeeping. The procedure will edit the default polling interval from 3,600 seconds (60 minutes) to 900 seconds (15 minutes) by editing the Windows registry. Then it will set the NTP servers from which you will synchronize time and change the behavior of the Windows time service from its default of following the domain hierarchy to using NTP servers. Which NTP servers you use to synchronize time from depends on your environment. Many organizations maintain stratum-1 NTP servers in-house, while others rely on external time servers such as those maintained by the United States Naval Observatory (`www.usno.navy.mil`).

Warning: Always take the proper precautions before editing the registry. Improper modification to the registry can leave your system in an unusable state.

1. Open the registry editor by opening the Start menu and typing **regedit.exe** in the Run or Search field.
2. Navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient.
3. Open or create the REG_DWORD value SpecialPollInterval.
4. The recommended value for the poll interval is 900 seconds, which is entered in base decimal.
5. Close the registry editor, and open a command prompt by opening the Start menu and typing **cmd** in the Run or Search field.

6. Enter the command `w32tm /config "/manualpeerlist:time1.mydomain.com,0x1 time2.mydomain.com,0x1"`.
7. Enter the command `w32tm /config /syncfromflags:MANUAL`.
8. For these settings to take effect, the Windows time service must be restarted. Enter the following command to restart the service: `net stop w32time && net start w32time`.
9. To force a synchronization of the time, enter the following command: `w32tm /resync`.
10. Close the command window by entering the following command: `exit`.

Using Templates to Deploy Virtual Machines

One of our favorite features of working with virtual machines is the portability and what that allows. It's no secret that one of the least desirable tasks of the system administrator is new system deployment. Not only do you find yourself unpacking boxes of memory, processors, and disks, but then you have to assemble, rack, cable, and, worst of all, go through the time-consuming process of OS installation. Those are precious hours that you spend clicking through ever-familiar dialog boxes, many Next and OK buttons, looking for the latest antivirus and backup software, and installing all of the updates to get the system ready for handoff.

Virtual machines have changed the way we deploy new systems. These days, vSphere administrators can deploy a new virtual machine in minutes. Using templates, vSphere administrators are able to create a base image that has been prepackaged with an OS and updates, software packages, and settings necessary for deployment. This is referred to as a *golden image*. There may be more than one image, but this is the baseline for all images. With vSphere, this golden image is converted from a virtual machine into a template. These templates can then be used to quickly deploy virtual machines that follow a consistent set of base configurations.

Templates appear alongside virtual machines in the vCenter client. They are stored in datastores and are made of the same files that make up the virtual machines, with one exception: The virtual machine will have a VMX file where the virtual machine configuration is kept. A template will have a VMTX file, which is the VMX file renamed to let vCenter and ESX know it is a template. We'll now show how to convert the golden image to a template.

Once the golden image virtual machine has been shut down, right-click the virtual machine in the vSphere client to bring up the context menu (Figure 3.25). In the virtual machine context menu under Template, you'll see two choices: Clone To Template or Convert To Template.

FIGURE 3.25
Right-clicking a virtual machine to convert it to a template by either cloning or converting the existing VM object

Clone To Template vs. Convert To Template

Clone To Template will make an exact copy of the virtual machine that will be labeled as a template. Changes made to the virtual machine will not be reflected in the cloned template. Cloning the virtual machine allows you to keep the golden image intact and use it as a base for other VM templates.

Convert To Template changes the existing virtual machine into a template. No copies or clones are made when converting to template. Once converted to a template, it can easily be converted back into a virtual machine to be powered on and updated.

Selecting Clone To Template launches a wizard similar to the New Virtual Machine Wizard. The wizard will walk you through selecting a name and location for the template, a host and data store to house the virtual machine/template files, and the disk format (same as source, thin, or thick disk provisioning). Click Finish on the summary screen to complete the wizard; this will initiate the clone operation. Once that's completed, you will have the original virtual machine as well as an exact replica labeled as a template within vSphere.

When cloning an existing virtual machine to a template, the template will be kept in the same container or folder as the virtual machine. You can move the template to a different folder or container just like the virtual machines.

You will notice that the icon changes from the standard virtual machine icon to the template icon (Figure 3.26).

Figure 3.26
Templates will not appear alongside virtual machines in the Host and Clusters view; you must switch to the VM And Templates view to view templates.

In addition to the icon change, the context menu (displayed when right-clicking the virtual machine/template) will be adjusted to display the available template options (Figure 3.27).

Figure 3.27
Templates have fewer options than virtual machines; they cannot be powered on, and you cannot edit their hardware.

After performing either of the two available actions, you will have a template ready for deployment or further customization. To deploy a new virtual machine, right-click the template, and select Deploy Virtual Machine From This Template from the context menu. This will launch a wizard similar to the New Virtual Machine Wizard.

When you've completed the wizard, the template will be cloned into a new virtual machine. At this point, the new virtual machine is now unique, but the OS will simply be an exact duplication of the source. In the next section, you will create a guest customization that will use the Sysprep features of Windows during virtual machine deployment to create a unique Windows installation while maintaining the configuration and software you built into the golden template.

Readying the Guest OS for Deployment

The process by which you prepare a guest operating system deployed from a template into a unique and usable virtual machine that you can deploy in your environment is called *guest customization*. These customization profiles are created using the vCenter client and are stored in the vCenter database. Windows administrators will notice the similarity between guest customization and Sysprep. Under the covers Sysprep is really doing the work, but vSphere gives you the user interface through vCenter to create the specifications using the customization wizard. This allows you to attach these customizations to templates running supported guest operating systems during deployment.

Guest customization allows you to preconfigure the following settings:

- Registration information
- Computer name
- Licensing information
- Administrator password
- Time zone
- Run-once commands
- Network settings
- Domain or workgroup membership
- New SID generation

The following steps will walk you through the process of creating a customization specification:

1. Launch the vSphere client, and click Customization Specifications Manager, located on the vSphere client Home page in the Management section (Figure 3.28).
2. Click New (Figure 3.29) to launch the vSphere Client Windows Guest Customization Wizard.
3. On the New Customization Specification page of the wizard, select the target OS as Windows (Figure 3.30), enter a descriptive name and a description if you like, and click Next.

FIGURE 3.28
You can find the Customization Specifications Manager on the Home page's Management section of the vCenter client.

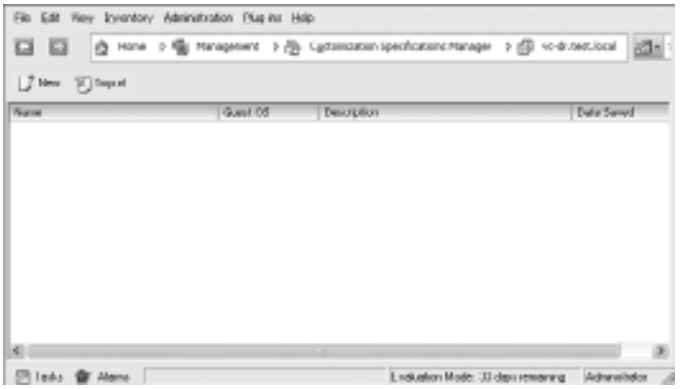

FIGURE 3.29
The Customization Specifications Manager displays all the available customization specifications and allows importing of previously exported specs.

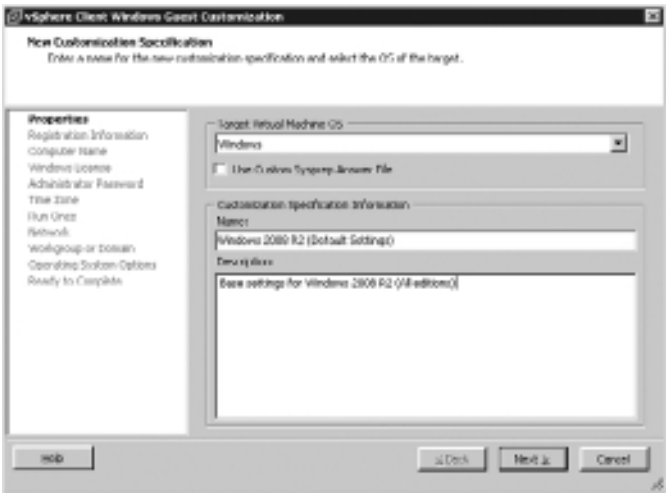

FIGURE 3.30
Selecting the target virtual machine guest OS and entering a descriptive name and detailed description

4. The name and organization will be displayed in Windows as the registered owner of the system. Enter the defaults for your organization, and click Next.

5. The computer name window gives some flexibility on how the guest OS computer name is assigned during Sysprep. The choices are as follows:

   Enter A Name: Manually entered names are OK for deploying one or two virtual machines but mostly not applicable to deployment with templates.

   Use The Virtual Machine Name: Using this option is the best practice and will ensure that the virtual machine name is consistent with the guest OS computer name. This makes it easy to find virtual machines using a single name.

   Enter A Name In The Deploy Wizard: This is similar to manually entering the name but geared more toward template deployments. The name is selected and applied during the deployment process.

   Generate A Name Using The Custom Application Configured With The vCenter Server: If your deployment process already employs an application that will autogenerate a computer name, it can be used here.

6. To keep with best practices, you will keep the names of the virtual machines and the guest OS in sync (Figure 3.31). Select the Use The Virtual Machine Name radio button, and click Next.

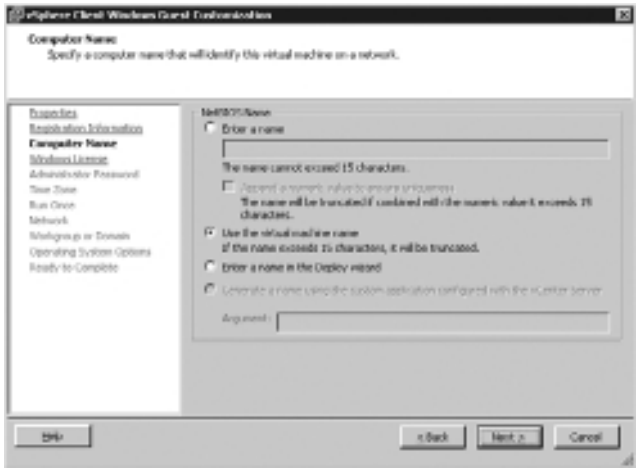

FIGURE 3.31
Using the virtual machine name is usually the best option to keep the names consistent and avoid confusion.

7. On the Windows License page, you can enter a product key. If the golden image was deployed using volume license media, you may leave the product key blank. Select the appropriate server license mode, Per Seat or Per Server, and click Next.

8. Enter an administrator password, which will be applied to the virtual machine guest OS during deployment. The administrator password screen also allows you to specify whether you want the local administrator account to log on and how many times.

9. After selecting the appropriate time zone, the Run Once screen is where you can specify an application or script to be run upon user login (Figure 3.32). If you want the application or script to run immediately without user intervention, you must select the option to have the local administrator account log in at least once on the administrator password screen.

FIGURE 3.32
Use Run Once when a custom script or application should be run after deployment and on first login.

10. The network settings will be different for each virtual machine. On the Network page, you can leave the default Typical Settings, which enables DHCP on all interfaces, or customize the settings. The template will use the default of typical settings, and you will customize the network settings during deployment. Select the Typical Settings radio button, and click Next.

11. Once deployed, the guest OS may be joined to an Active Directory domain or a workgroup (Figure 3.33). To join an Active Directory domain, you must enter the credentials of a user account with the permission to add the computer account to the domain. These credentials will be saved in the customization specification and used during deployment of the template.

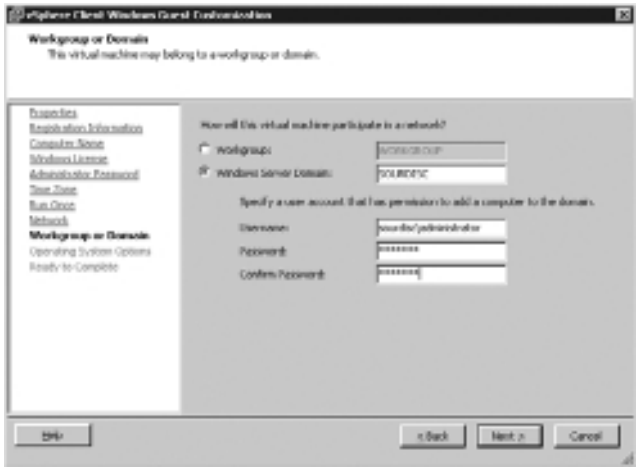

FIGURE 3.33
In an Active Directory domain environment, having a dedicated account to add deployed virtual machines to the domain can help reduce administrative overhead and auditing.

12. Select the box to generate a new machine security ID (SID), and click Next.
13. Review the settings, and click Finish to complete the new guest customization specification (Figure 3.34).

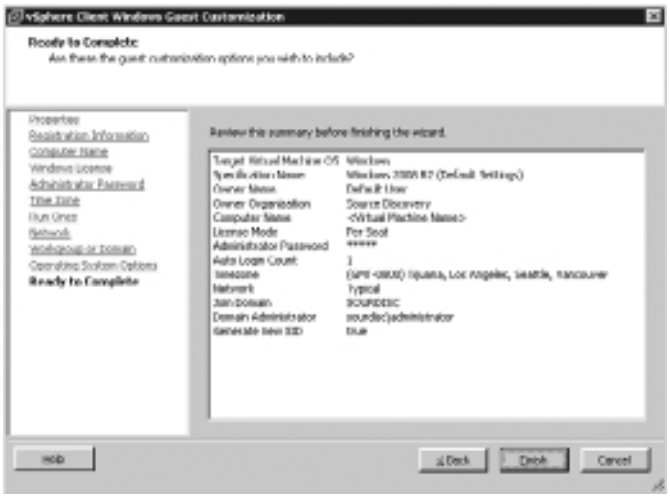

FIGURE 3.34
Review the customization summary to ensure settings are correct.

Once complete, you can use the Customization Specifications Manager to work with the saved customizations. Customizations can be cloned, renamed, edited, and deleted.

DEPLOYING VIRTUAL MACHINES FROM A TEMPLATE

With your template and a customization specification created, you can begin deploying virtual machines that will be configured to your specifications, will be unique, and will be functional in minutes as opposed to hours. To begin, follow these steps:

1. Right-click the template, and select Deploy Virtual Machine From This Template in the context menu. The first five screens will look familiar because they are the same as you saw when deploying new virtual machines. These screens allow you to name (which will also name the guest OS) and select a location for your VM, select a compatible host or cluster, place it in a resource pool if desired, and select the provisioning of the virtual disk and where to place the VM files.

   The Guest Customization page (Figure 3.35) gives you a few choices on what to do during deployment:

   Do Not Customize: Clone the VM or template without applying any new settings to the guest OS.

   Customize Using The Customization Wizard: Use the customization wizard to manually set all customizable settings. The data entered here will be used during the Sysprep process. This includes changing the IP address, name, licensing information, and SID.

Customize Using An Existing Customization Specification: Use a preconfigured customization specification to set all of the predefined fields during the Sysprep process.

Use The Customization Wizard To Temporarily Adjust The Specification Before Deployment: This is an add-on to the previous option and allows you to predefine all of the generic settings such as organization, computer name generation, licensing, and so on, while still providing the flexibility to set guest OS–specific information such as IP address and domain membership.

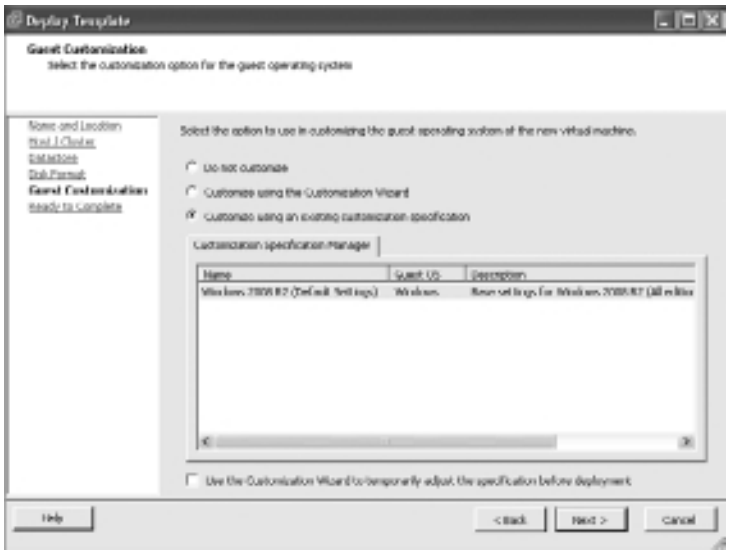

FIGURE 3.35 Guest Customization page during template deployment

2. Select the Customize Using An Existing Customization Specification radio button, and choose the specification you created earlier, in other words, Windows 2008 R2 (Default Settings). Then click Next.

TIP For the greatest level of flexibility and ease of management, it's best to create the customization specifications using as little guest OS–specific information as possible. This includes computer names and IP addresses. Standardize on information such as registration information, licensing, and domain membership, and select the Use The Customization Wizard To Temporarily Adjust The Specification Before Deployment box. This feature allows you to use the prepopulated information in your existing specification as well as make changes to specific portions that are relevant to the new VM and guest OS.

On the Ready To Complete page, you can review the settings that your new virtual machine will use, verify the customization specification is correct, and power on the VM after creation or (experimentally) go directly into editing the hardware before creation.

3. Check the Power On This Virtual Machine After Creation box, and click Finish (Figure 3.36).

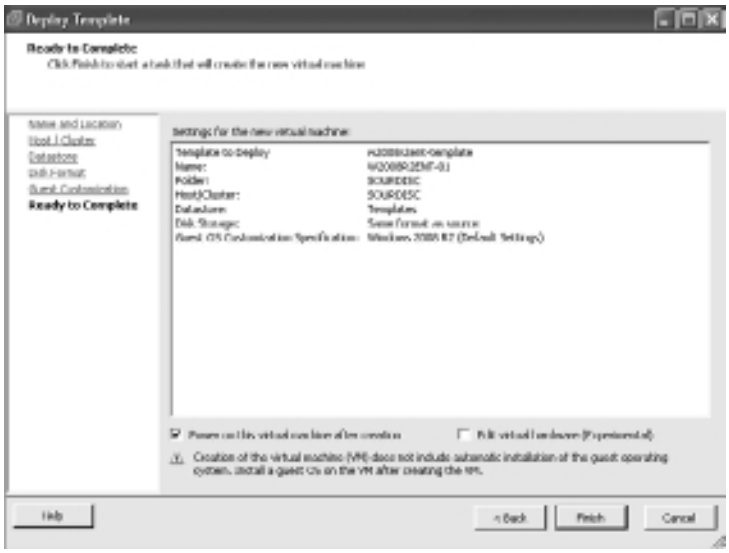

FIGURE 3.36
Check the Power On This Virtual Machine After Creation box to allow Sysprep to complete

TIP Allow the deployment wizard to power the VM on after creation. When choosing to customize a guest OS, you should be aware that the customization does not take place until after the deployed VM has been powered on. The first boot is required to place the customization files into the VM. After the files are present, the VM is rebooted, which kicks off the Sysprep process and then rebooted once more after completion to commit all changes. For this reason, it's generally a good idea to select the Power On This Virtual Machine After Creation box and let the deployment wizard do its thing.

WARNING Don't hit Ctrl+Alt+Del just yet. If you are sitting at the console waiting to log into your fresh Windows 2008 R2 VM, sit tight when you see the Ctrl+Alt+Del screen for the first time. The Sysprep process is running in the background, and chances are you will get kicked out of your session while the system restarts for the last time.

After a few minutes, you will be presented with the logon screen to your new Windows 2008 R2 VM and be ready to move on to deploying any additional features or roles.

Converting a Physical Machine to a Virtual Machine

Up to now we've concentrated on deploying virtual machines by creating a new VM or deploying from a template. However, chances are that you have a few physical machines that are newly built or that have highly customized applications that you can't (or don't want to) rebuild completely as virtual machines. Well, you're in luck.

VMware vCenter Converter allows you to convert a physical machine (running or powered off) to a virtual machine. On a running Windows machine, this is done by installing an agent remotely onto the source physical machine. The agent takes a VSS snapshot of the source volume and creates a virtual machine on the destination ESX host. The VSS snapshot is used to populate the virtual disk file. The agent installs the required drivers in the virtual machine to

allow the VM to boot successfully using the new virtual hardware and then removes itself from the source physical machine.

During a cold clone, the source physical machine is rebooted into the VMware vCenter Converter boot CD. The boot CD contains its own operating system and the vCenter Converter application necessary for copying the volumes and creating the virtual machine on the destination ESX host. Once the virtual machine is created and the source volumes are copied to the virtual disk file, the virtual machine is updated with the appropriate drivers for the new virtual hardware. The source physical machine is left in the current state and can be rebooted into the original operating system or decommissioned. The virtual machine is ready for use.

VMware vCenter Converter is a plug-in module for vCenter Server, which can be used to perform the following:

- Converting existing physical machines into virtual machines and importing them into a vSphere environment
- Converting Microsoft Virtual Server VMs into VMware-format virtual machines
- Converting vCenter-managed virtual machines into formats usable by other VMware products, such as VMware Workstation and Server
- Restoring VMware Consolidated Backup images
- Converting third-party backup images into vCenter-managed virtual machines (such as Symantec Backup Exec and Norton Ghost images)

VMware vCenter Converter consists of the following components:

vCenter Converter Server Usually installed on the vCenter server but can also be installed on an independent server. This communicates to the vCenter client plug-in to import or export virtual machines.

vCenter Converter Client Deployed as a vCenter client plug-in and provides access to the vCenter Converter wizards.

vCenter Converter Agent Installed on a physical machine by the vCenter server to assist in the creation of the virtual machine; the agent is uninstalled automatically after the conversion process.

vCenter Converter Boot CD (optional) Allows for cold cloning of a physical machine.

Although this can be a very helpful way to quickly get your foot into the virtual door or migrate a physical host to reclaim hardware or rack space/power/cooling, watch out for the following:

Additional software Most hardware vendors install agents, dashboards, diagnostic tools, and so on, during deployment. Most of these tools become unnecessary when the physical host is migrated onto virtual hardware. Follow up after a conversion, and remove any unnecessary software from the guest OS.

Disk space Unless the server hardware is very old, chances are the OS has a few hundred gigabytes of storage space allocated to it while the system itself is using a fraction of the allocated space. When this is the case, you can reduce the disk size in the virtual machine, but this will require the migration to use a slower file-by-file cloning method as opposed to the block method that is the default.

You can find more information on VMware vCenter Converter at the vCenter Converter home page: `www.vmware.com/products/converter/`.

Deploying a Windows Failover Cluster in vSphere

As the proliferation of the virtual infrastructure continues, companies have virtualized the low-hanging fruit and are moving onto tier-1 applications. These are the mission-critical applications that your customers depend on being available upward of 99.9 percent of the time. How you provide that level of uptime really depends on the application being deployed. Many applications must rely on redundant hardware, DNS round-robin, hardware or software load balancing, or other application-independent methods of providing better availability.

vSphere provides the ability to configure ESX/ESXi hosts into a pool of resources called a *cluster*. The cluster manages the resources of all hosts as a single pool of CPU, memory, storage, and network. Clusters provide the VMware HA feature. VMware HA protects virtual machines by restarting them on an alternate host in the case of the original host failing. In most situations, VMware HA provides sufficient levels of availability for applications where a restart of the virtual machine is tolerable in the case of hardware failure.

If the restart of a virtual machine cannot be tolerated in the case of hardware failure, the Fault Tolerance (FT) feature of vSphere may be used. FT uses vLockstep technology to provide zero downtime and zero data loss for the virtual machine and its application in the case of a hardware failure. When a virtual machine is protected using FT, a second virtual machine runs in virtual lockstep with the primary. The second virtual machine runs isolated from the network on a second ESX/ESXi host and executes the same set of instructions as the primary virtual machine. This is done with no additional configuration to the virtual machine's guest operating system or applications.

With the options available using vSphere, why would you want to deploy a failover cluster between virtual machines? The fact is you are now deploying business-critical applications that in certain scenarios cannot be down whether because of hardware, software, or application outages or even monthly patching. Applications such as Exchange and SQL Server have utilized clustering features found in Microsoft failover clusters (formerly Microsoft Clustering Service) to mitigate the downtime required due to planned and unplanned outages. Luckily for those of us who run environments that demand this level of uptime, running virtual machines does not mean you can't use Microsoft Failover Clustering.

vSphere 4 supports Microsoft Cluster Service (MSCS) running on Windows 2000 and Windows 2003 and Microsoft Failover Clustering (FC) for Windows 2008. Support includes clustering on 32-bit and 64-bit guests and "Boot-from-SAN" virtual machines (boot/system virtual disk located on Fibre Channel–attached storage). Majority node-set clusters are also supported, and there is continued support for traditional shared storage clusters.

Clustering Virtual Machines

When deploying Microsoft Failover Clustering in a virtual environment, you can take a few different approaches:

Two virtual machines, one host A "cluster in a box" allows you to use a single ESX/ESXi host to build a cluster using two virtual machines housed on that host. This is great for testing or a lab environment, but it's not something you would want to run in production.

One virtual machine, one native machine In this scenario, the virtual machine is clustered with a native system. This method can be used to migrate an existing cluster from a native environment into a virtual environment.

Two virtual machines, multiple hosts A "cluster across boxes" will protect against software and hardware failures by placing the virtual machines on separate physical hosts. When deploying in a fully virtualized environment to production, this is the only recommended path and is what you will be focusing on next.

Hardware and Software Requirements for Failover Clusters

VMware requires specific hardware and software configurations to be followed when deploying failover clusters. Table 3.5 lists the virtual hardware requirements.

Table 3.5: Virtual Hardware Requirements

Component	Requirement
Operating system	Windows 2008, Windows 2008 R2
Virtual SCSI adapter	LSI Logic SAS (Windows 2008, Windows 2008 R2)
Shared storage type	Pass-through RDM (physical compatibility mode)
Virtual NIC	Default (e1000 or VMXNET3)
Virtual hardware version	7 (ESX/ESXi 4.0 and higher)

Along with the virtual hardware listed in Table 3.5, the following should be noted when deploying clustered virtual machines:

- The heartbeat network required for cluster communication should use a dedicated vSwitch and physical Ethernet adapter (vmnic). These vmnics can be connected to an unmanaged switch or to a private VLAN and must be present on all ESX/ESXi hosts where the clustered virtual machines will run.
- Virtual disks should be created using the Support Clustering Features Such As Fault Tolerance option. This is only for the nonshared disks in the VM to allow for the best performance possible. If you recall earlier, you used this feature when building your virtual machine golden image.
- Set the I/O timeout within the guest OS by modifying HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\TimeOutValue to a value of 60 seconds or greater. This will help avoid cluster failovers because of any delays in accessing storage.
- Shared disks will be set up on new SCSI controllers to allow enabling physical SCSI bus sharing.
- If virtual machine boot disks are located on a SAN-based VMFS volume, a Fibre Channel SAN must be used (iSCSI is not supported), and the StorPort LSI Logic drivers should be used in place of the SCSIport drivers.

Limitations of Failover Clustering in vSphere

The following bullets outline unsupported features when running Microsoft Failover Clustering within a vSphere environment. Only the virtual machines engaged in the failover cluster are affected by these limitations.

- Clustering on iSCSI or NFS disks
- Mixed environments where one vSphere cluster node is running on an ESX host and the second vSphere cluster node is running on an ESXi host
- Clustered virtual machines protected by VMware HA, DRS, or FT. These features can be disabled at the virtual machine level to allow other VMs in the cluster to take advantage of these vSphere features.
- Migration using VMotion of clustered virtual machines
- N-Port ID Virtualization (NPIV)
- ESX storage path policy set to round-robin when using native multipathing (NMP)

Building a Microsoft Failover Cluster

Now that you know what your supported configurations are and what to look out for, you can begin building your failover cluster. In this section, you will build a two-node cluster using shared storage. Table 3.6 shows the virtual machine configuration you will follow, and Table 3.7 shows the disk configuration for the cluster.

Table 3.6: Cluster Node Build Specifications

Component	Details
Virtual machines	
Node1 VM/hostname	clsqlsrv01
Node2 VM/hostname	clsqlsrv02
CPU/memory	
CPU	4
Memory	8192 MB
Network	
Production network interface	VM Network
Heartbeat network interface	HeartBeat

TABLE 3.7: Shared Disks

DISK USE	DRIVE LETTER	BUS:DEVICEID	SIZE
OS	C:	0:0 (nonshared)	40 GB
Witness disk	W:	1:0 (Shared)	1 GB
Shared disk	S:	1:1 (Shared)	20 GB

Here's a quick overview of the tasks you will be performing:

1. Deploy two virtual machines from a template.
2. Add network, SCSI controllers, and disks to node 1.
3. Add network, SCSI controllers, and disks to node 2.
4. Install and configure cluster services.
5. Test failover.

CLUSTER NODE 1 CONFIGURATION

In the next couple of exercises, you will deploy the two virtual machines that will make up your cluster. Although they are mostly identical, you will notice slight differences to the way the disks are configured for each virtual machine. This is crucial to the functionality of the cluster and the shared disks. You'll begin with configuring cluster node 1:

1. Deploy two virtual machines as described in the previous section using the specifications in Tables 3.6 and 3.7.
2. Once you've completed the virtual machines, the Sysprep process on both can be shut down.
3. Open the virtual machine properties for clsqlsrv01.
4. Click Add to add new virtual hardware.
5. In the Device Type window, select Ethernet Adapter.
6. In the Network Type window, select the adapter type as VMXNET3. Under Network Connection, select the port group dedicated to cluster private communications, HeartBeat in this case. The Connect At Power On box should be selected.
7. Click Finish on the Ready To Complete page to complete the wizard (Figure 3.37).
8. To begin adding the disks, click Add in the virtual hardware properties.
9. Select Hard Disk, and click Next.

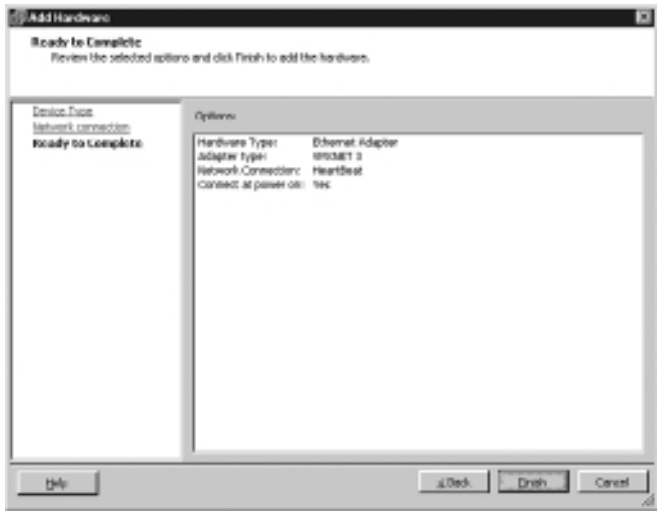

FIGURE 3.37
Click Finish to commit the new Ethernet adapter.

10. On the Select A Disk page, select Raw Device Mappings, and click Next (Figure 3.38).

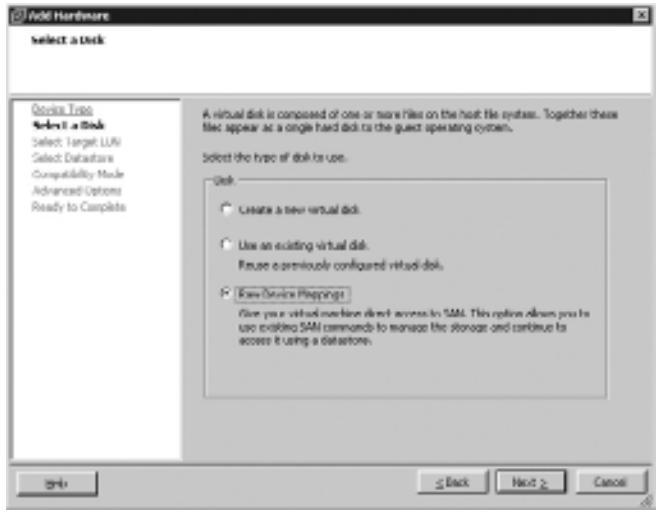
Add Hardware
Select a Disk
Device Type
Select a Disk
Select Target LUN
Select Datastore
Compatibility Mode
Advanced Options
Ready to Complete
A virtual disk is composed of one or more files on the host file system. Together these files appear as a single hard disk to the guest operating system.
Select the type of disk to use.
Disk
Create a new virtual disk
Use an existing virtual disk
Reuse a previously configured virtual disk.
Raw Device Mappings
Give your virtual machine direct access to SAN. This option allows you to use existing SAN commands to manage the storage and continue to access it using a datastore.
Help
< Back
Next >
Cancel

FIGURE 3.38
Raw device mappings are required for Microsoft Failover Clusters across hosts.

11. The Select And Configure A Raw LUN page displays the available, unformatted LUNs available for allocating as raw device mappings (Figure 3.39). Select the first shared disk (1GB Witness disk in this configuration), and click Next.

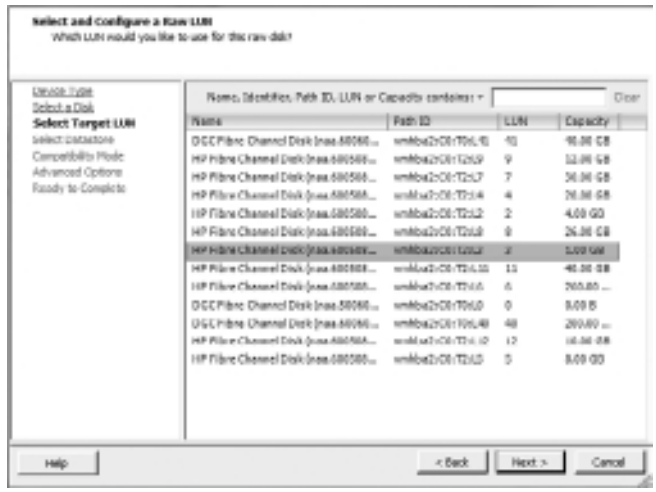

FIGURE 3.39
The Select And Configure A Raw LUN page will display only those LUNs that are non-VMFS data stores and available for use as a raw device.

12. When a raw device mapping is used, a small VMDK file is created on a VMFS data store (Figure 3.40). This VMDK file is a pointer to the RDM and will be used when you configure the second node in the cluster. Select a VMFS data store in which to house the RDM pointer files, and click Next.

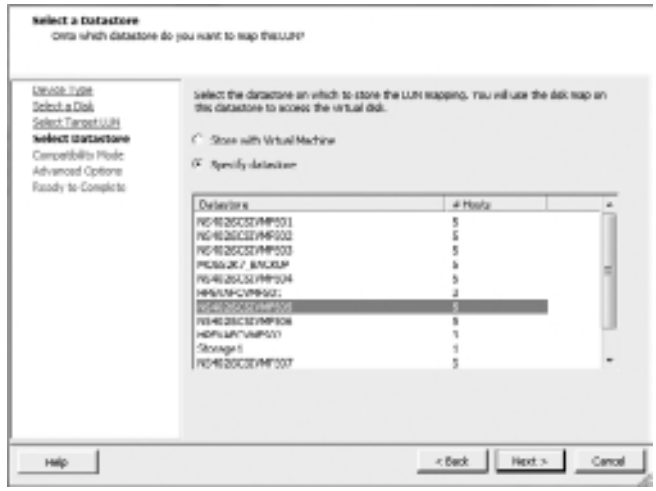

FIGURE 3.40
Specify which data store will hold the raw device pointer files.

13. On the Select Compatibility Mode page, select the Physical radio button (Figure 3.41), and click Next.

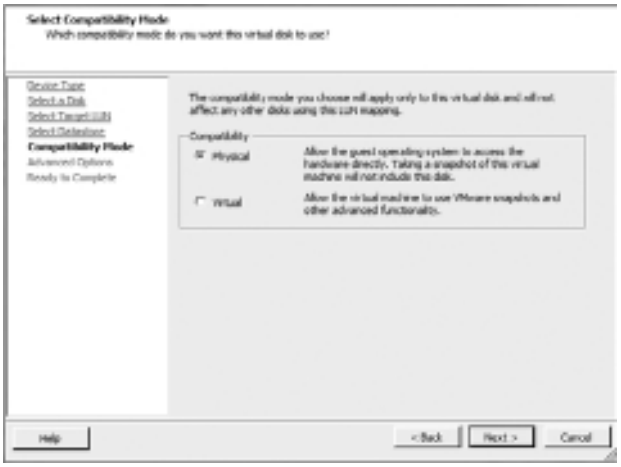
Select Compatibility Mode
Which compatibility mode do you want this virtual disk to use?
Device Type
Select a Disk
Select Target LUN
Select Datastore
Compatibility Mode
Advanced Options
Ready to Complete
The compatibility mode you choose will apply only to this virtual disk and will not affect any other disks using this LUN mapping.
Compatibility
Physical
Allow the guest operating system to access the hardware directly. Taking a snapshot of the virtual machine will not include this disk.
Virtual
Allow the virtual machine to use VMware snapshots and other advanced functionality.
Help
< Back
Next >
Cancel

FIGURE 3.41
The physical compatibility mode is required for Microsoft Failover Clusters across hosts.

NOTE *Disk compatibility mode* describes the method by which the virtual machine accesses the underlying storage. Physical compatibility mode, required when deploying shared-storage clusters, allows the virtual machine to directly access the storage. This method allows for array-based snapshots or clones of the physical disks when using agents within the guest OS. Snapshots of the virtual machine will not include disks in physical compatibility mode. Virtual compatibility gives you many of the benefits of physical compatibility mode while adding just enough virtualization to support features such as virtual machine snapshots.

14. The Advanced Options page is where you will select the SCSI bus and device target number for your RDM. Because this is a shared disk, you must install this disk on a new SCSI controller. You can do this by selecting the first device on a new bus, that is, 1:0, 2:0, 3:0, and so on (Figure 3.42). Select SCSI (1:0), and click Next.

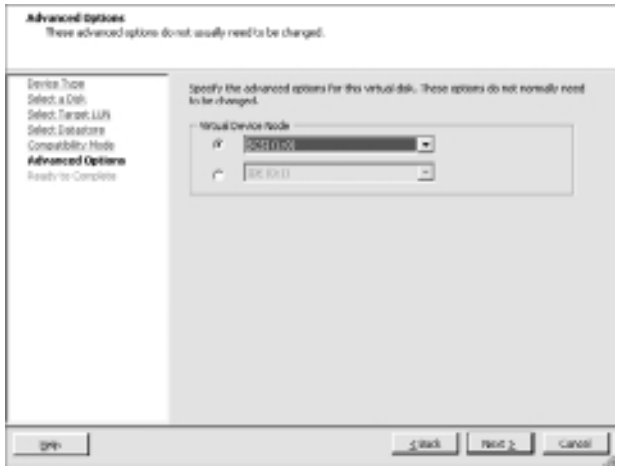

FIGURE 3.42
Select a new SCSI bus, that is, 1:0 or 2:0, to create a new SCSI controller to be shared between hosts.

15. On the Ready To Complete page (Figure 3.43), you can review the new disk configuration and click Finish.

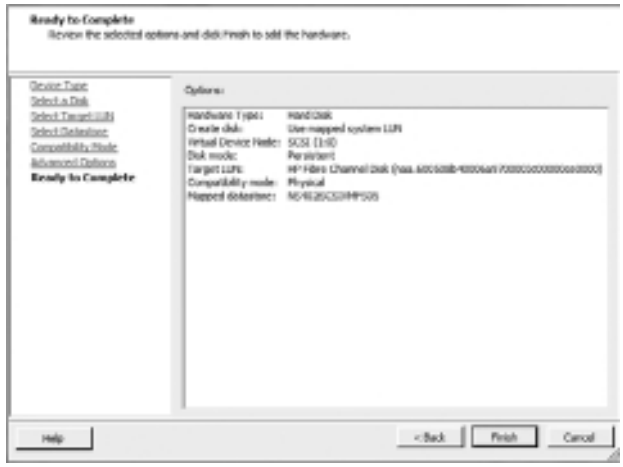

FIGURE 3.43
Review the details of the new raw device.

TIP Because you added the new disk as a new device on a new bus, you now have a new virtual SCSI controller. The SCSI controller will default to the recommended type, LSI Logic SAS for Windows 2008 and Windows 2008 R2.

16. To support your shared storage cluster, you must allow SCSI bus sharing on the new controller. In the virtual machine properties, select the new SCSI controller, and select the Physical radio button on the right, under SCSI Bus Sharing (Figure 3.44).

17. Repeat steps 8–16 for the remaining disks.

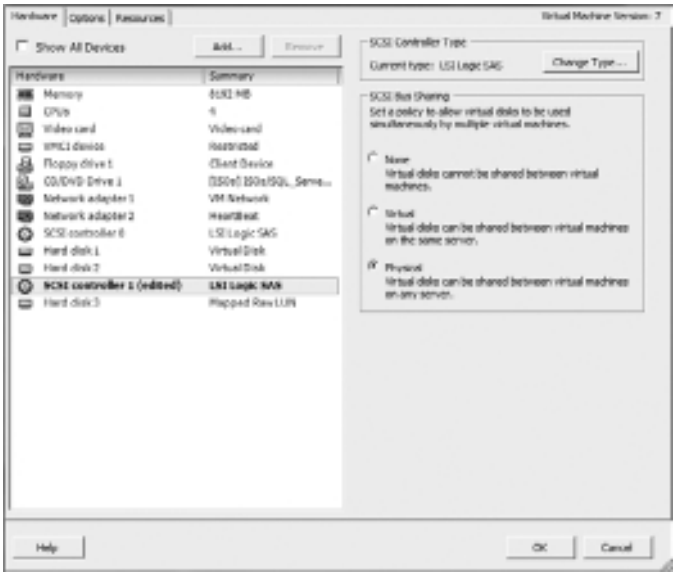

FIGURE 3.44
Physical SCSI bus sharing is required for failover clustering to detect the storage as usable.

WARNING SCSI Controller 0 must not be enabled for SCSI bus sharing. Not only is this an unsupported configuration, but it may result in data corruption because the disk(s) would not be under the control of the clustering service. Any disks that will be part of a cluster group need to be located on a new SCSI controller with physical bus sharing enabled.

Cluster Node 2 Configuration

With the first node of the cluster complete, you can begin configuring the second node. In the configuration of the first node, you defined all of the shared disks as raw device mappings and created the associated VMDK pointer file. You'll notice during the configuration for the second node that you will simply point to the VMDK pointer file for the shared disks instead of attempting to create new raw device mappings. vSphere will recognize these as being pointers for raw device mappings and make the appropriate adjustments.

1. Open the virtual machine properties for clsqlsrv02.
2. Click Add to add new virtual hardware.
3. On the Device Type page, select Ethernet Adapter.
4. On the Network Type page, select the adapter type as VMXNET3. Under Network Connection, select the port group dedicated to cluster private communications, HeartBeat in this case. The Connect At Power On box should be selected.
5. Click Finish on the Ready To Complete page to complete the wizard.
6. To begin adding the disks, click Add in the virtual hardware properties.
7. Select Hard Disk, and click Next.
8. Because you have already assigned the raw LUN to the first node in the cluster, you can now point the second node to the RDM pointer file (Figure 3.45). To do this, select Use An Existing Virtual Disk, and click Next.

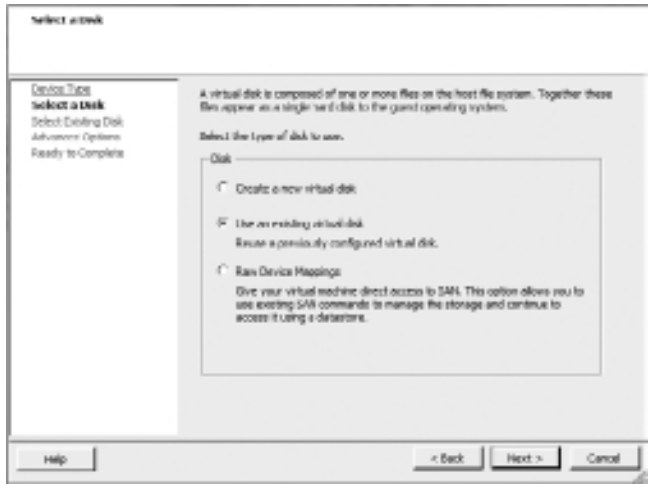

Figure 3.45
Since the RDM pointer file has already been created, you only need to point to that already existing VMDK file.

9. To select the RDM pointer file, click the Browse button, and navigate to the data store (Figure 3.46) and folder (clsqlsrv01) where you have created the pointer file (Figure 3.47). Select the pointer file for the first disk (`CLSQLSRV01.vmdk`), and click OK and Next to proceed.

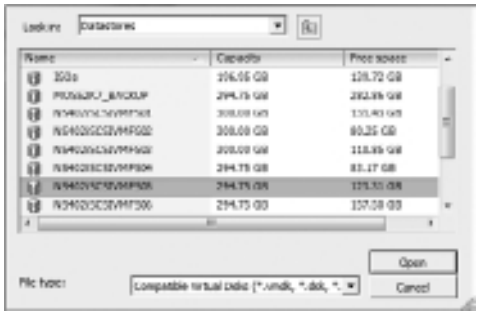

FIGURE 3.46
Select the data store that was used to store the RDM pointer file when created on cluster node 1.

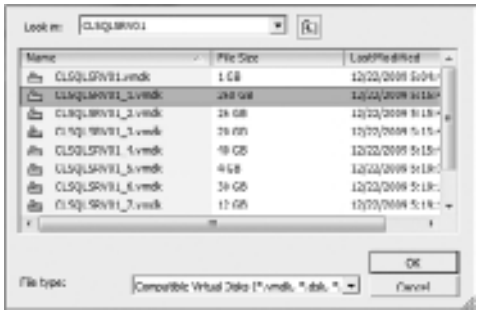

FIGURE 3.47
Selecting the pointer file

10. Select the SCSI bus and device target number based on Table 3.7. These will match the bus and device target numbers used on the first cluster node. Click Next to continue.
11. Review the options on the Ready To Complete page, paying special attention to the Virtual Device Node and Disk File Path options. Click Finish once you've verified all settings.
12. As with the first cluster node, you will have to set the SCSI Bus Sharing option to Physical and verify the type is LSI Logic SAS.
13. Repeat steps 5–11 for the remaining disks. When completed, the virtual hardware inventory will look similar to Figure 3.44.

Failover Cluster Configuration

With both cluster nodes configured from a virtual hardware perspective, you can begin configuring the cluster from within the guest.

1. Power on cluster node 1 (clsqlsrv01).
2. Open a console connection to node 1, and log into the guest.

3. Launch Server Manager, and navigate to Disk Manager under Storage.
4. In Disk Manager, you will see the new disks as being offline. Right-click each disk label, and select Online (Figure 3.48). If necessary, right-click the disk label again, and select Initialize Disk (Figure 3.49). Then select the MBR partition type (Figure 3.50).

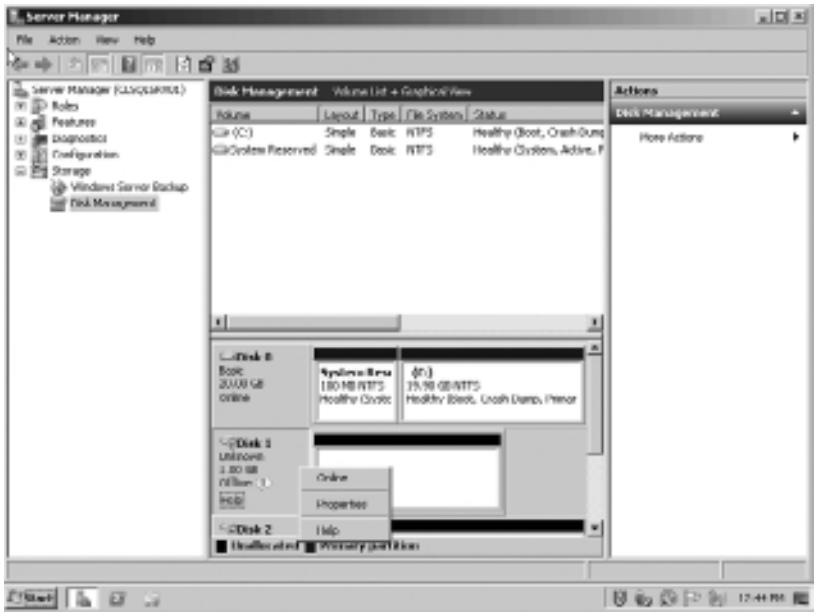

FIGURE 3.48 Bringing the disk online to be able to initialize

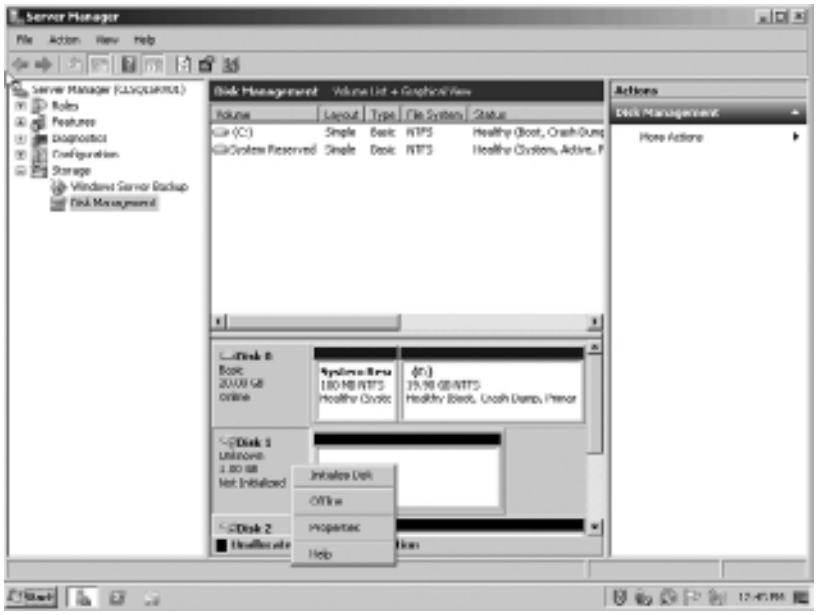

FIGURE 3.49 Initializing the disk to allow Windows to create a new partition

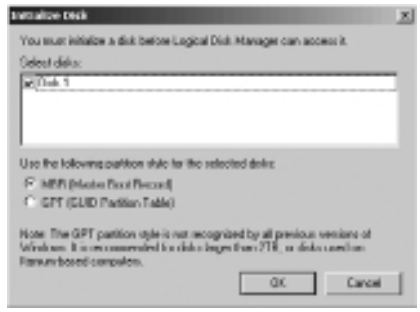
Initialize Disk
You must initialize a disk before Logical Disk Manager can access it.
Select disks:
Disk 1
Use the following partition style for the selected disks:
MBR (Master Boot Record)
GPT (GUID Partition Table)
Note: The GPT partition style is not recognized by all previous versions of Windows. It is not supported for disks larger than 2TB, or disks used on Itanium-based computers.
OK
Cancel

FIGURE 3.50
Select the disk or disks if multiple ones were added, and click OK.

5. Once the disk appears as online, right-click the unallocated volume, and select New Simple Volume.

6. Follow the New Simple Volume Wizard to define the size and drive letter. For data disks such as those holding databases, the allocation unit size should be set to 64K, and a full format is recommended. See Figure 3.51 for the result.

New Simple Volume Wizard
Completing the New Simple Volume Wizard
You have successfully completed the New Simple Volume Wizard.
You selected the following settings:
Volume type: Simple Volume
Disk selected: Disk 2
Volume size: 20477 MB
Drive letter or path: S
File system: NTFS
Allocation unit size: 65536
Volume label: Shared Data
To close this wizard, click Finish.
< Back
Finish
Cancel

FIGURE 3.51
Reviewing the settings on the summary page of the New Simple Volume Wizard

Disk Alignment

Disk alignment is still a factor to consider when deploying high-performing disks within a virtual machine. In a standard MBR disk, the first 63 sectors (starting at 0) are usually reserved or hidden from use. Windows versions prior to Windows 2008 create the partition on the first available sector (sector 63), which subsequently causes some writes (depending on file allocation unit and stripe unit size) to cross the physical boundaries. This in turn can cause an increased number of I/Os, which can lead to performance degradation.

In versions of Windows prior to Windows 2008, it was considered a best practice to align the partitions manually at the time of creation using `diskpar.exe` or `diskpart.exe`. Windows 2008 uses a default starting offset of 1 MB (1,048,576 bytes) and in many cases is the recommended setting; however, consult with your storage vendor or SAN administrator to make sure these settings are suitable for your environment.

7. With all disks online, initialized, and formatted on node 1, power on node 2 (clsqlsrv02).
8. Log into node 2, and open Disk Management from within Server Manager.
9. Right-click each disk label, and select Online; once the disks are online, you will see the volume labels and status.
10. The disks have been assigned the next available drive letters. You'll need to change the drive letters to match the letters you've assigned on node 1 (refer to Table 3.7). Right-click the partition on the disk labeled Witness, select Change Drive Letter And Path, click Change, assign the drive letter W, and click OK. You will be warned about dependencies of the drive letter. You have no data on these drives yet, so you may safely click Yes and continue.
11. Repeat the process for the disk labeled Shared, this time assigning drive letter S and again ignoring the warning because you have yet to populate the drive with any data (Figure 3.52).

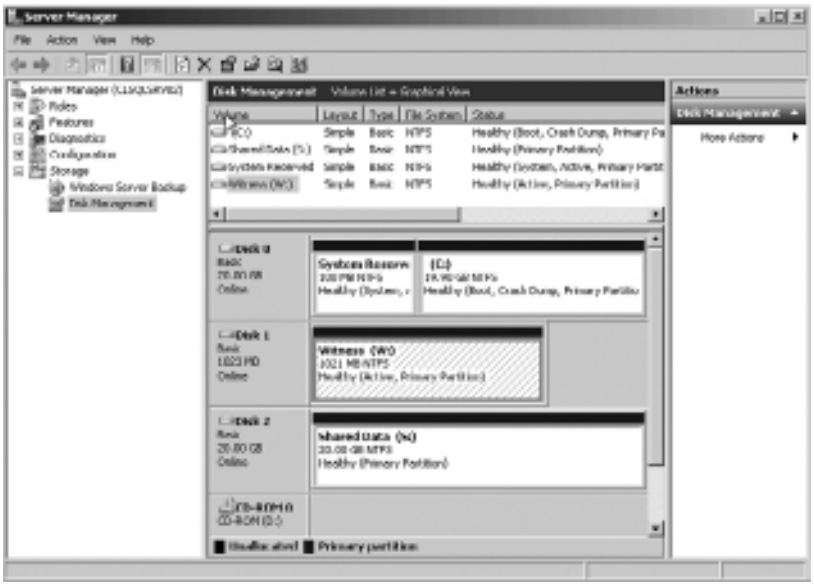

FIGURE 3.52 CLSQLSRV02 will look similar to CLSQLSRV01 when completed.

12. On each node of the cluster, follow steps 13–16 to install the Failover Clustering feature.
13. From within Server Manager, under Features, click Add Features.
14. Select Failover Clustering, and click Next (Figure 3.53).
15. Review the information in the confirmation window, and click Install.
16. When the install is complete, you can review the results (Figure 3.54) and click Close to complete the wizard.

FIGURE 3.53
Select the Failover Clustering feature for installation.

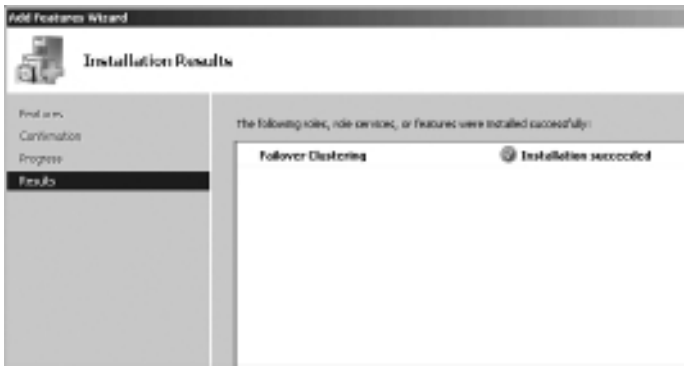

FIGURE 3.54
Failover Clustering feature successfully installed

17. In Server Manager on node 1, click Failover Cluster Manager under Features; this will display the tasks you can perform.
18. Under the Management section, click Validate A Configuration.

NOTE Validation will run a variety of tests against your (virtual) hardware, including the storage and networking, to verify whether the hardware is configured correctly to support a failover cluster. To pass all of the tests, both nodes must be online, and the hardware must be configured per the previous steps.

19. On the Before You Begin page, you can read more about the validation process and tests that will be performed; then click Next.
20. Enter the computer name of each node in the cluster, and click Add; the computer names will be validated, and each one will be placed in the Selected Servers field. Click Next to continue.
21. Select Run All Tests under Testing Options, and click Next.
22. On the confirmation screen (Figure 3.55), you may review the tests that will be performed and click Next to continue.

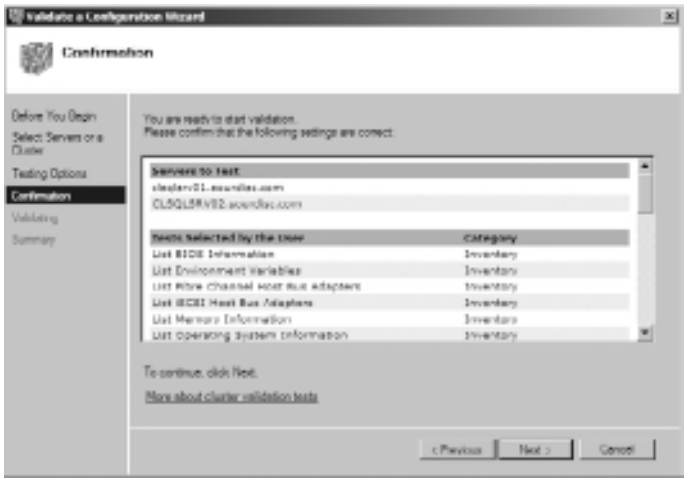

FIGURE 3.55
Run the validation wizard to ensure the cluster nodes are set up correctly.

23. Once the tests have completed, click the View Report button to see a detailed explanation of all of the tests and the outcome (Figure 3.56). On the Summary page, you can also kick off the cluster creation process. Click the Create The Cluster Now Using The Validated Nodes link to continue.

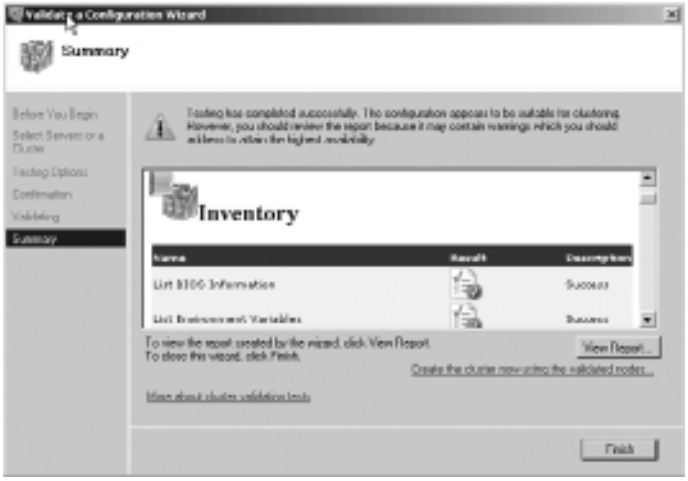

FIGURE 3.56
Verify there are no critical errors; the warning here indicates a lack of Internet connectivity.

24. The Create Cluster Wizard kicks off with the standard Before You Begin page; click Next to continue.

25. The access point is the name and IP address you will use to administer or access clustered resources (Figure 3.57). If you recall, your cluster nodes have two networks, one for private cluster communication (HeartBeat) and the other for access from the general LAN (VM Network). The IP address you provide will be on the network labeled VM Network. Enter a name and an available IP address, and click Next.

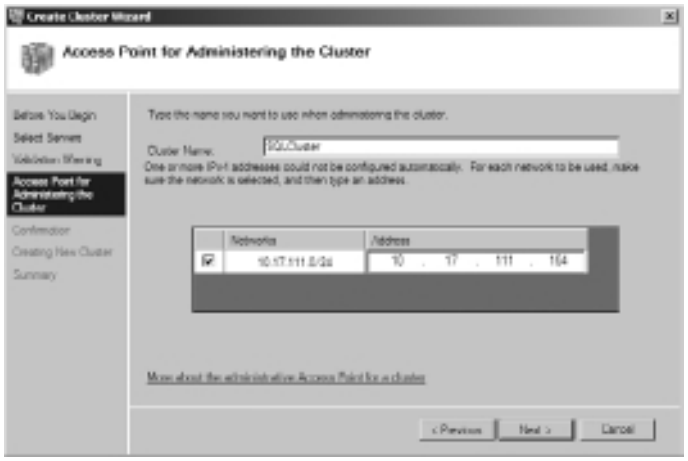
Create Cluster Wizard
Access Point for Administering the Cluster
Before You Begin
Select Servers
Validation Warning
Access Point for Administering the Cluster
Confirmation
Creating New Cluster
Summary
Type the name you want to use when administering the cluster.
Cluster Name:
SQLCluster
One or more IPv4 addresses could not be configured automatically. For each network to be used, make sure the network is selected, and then type an address.
Networks
Address
10.17.111.0/24
10 . 17 . 111 . 164
More about the administrative Access Point for a cluster
< Previous
Next >
Cancel

FIGURE 3.57
Enter the IP address that will be used to access the cluster resources.

26. On the next page, confirm the settings are correct, and click Next (Figure 3.58).

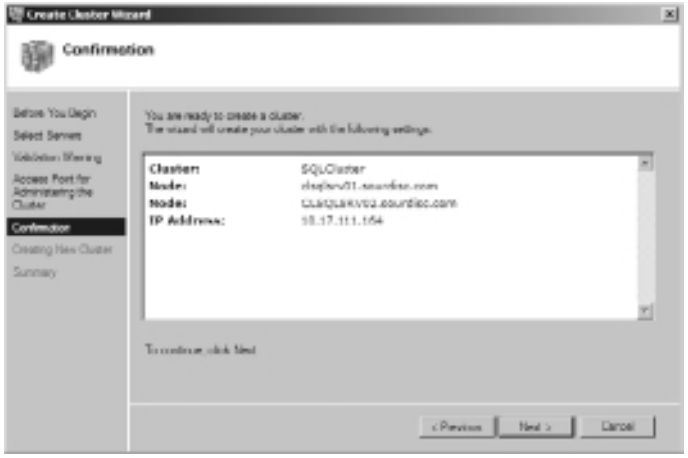

FIGURE 3.58
Verify all cluster details are correct, and click Next to begin creating the cluster.

27. On the Summary page, review the results (Figure 3.59), and click Finish.

FIGURE 3.59
The Summary page confirms the creation of the cluster and gives you the option to view the report, which also can be exported and saved.

Congratulations, you have now configured a Windows 2008 failover cluster across physical hosts using virtual machines! We will refer to this cluster in later chapters, which will allow you to complete an enterprise-ready deployment of Microsoft SharePoint using clustered Microsoft SQL Server databases.

Protecting Your Data

The topic of data protection is one that makes people cringe even though we all have to endure the task of having a solid backup or, more important, a solid recovery solution. The fact that you're moving everything off physical hardware and into, basically, files on disk would make it appear as though backups are as easy as just copying those files from disk to some other media and shipping it off-site for a period of time. Although that may somewhat be the case with some available solutions, you still have applications running within these virtual machines, and that makes it a bit more difficult than just copying files from a VMFS data store. In the following sections, you'll learn about your options for backing up the data within the virtual machine first, and then you'll explore your options for backing up the entire VM.

In-Guest Backup Solutions

For smaller environments, the Windows Server Backup feature is able to meet many of the requirements of a backup/restore solution. Besides being available in all versions of Windows 2008, the Windows Server Backup (Figure 3.60) feature provides the following:

- Volume Shadow Copy Service (VSS) and block-level backup technology
- Wizard-based backup and restore operations
- Command-line interfaces using `wbadmin.exe` and PowerShell cmdlets
- Full OS recovery using the Windows Recovery Environment

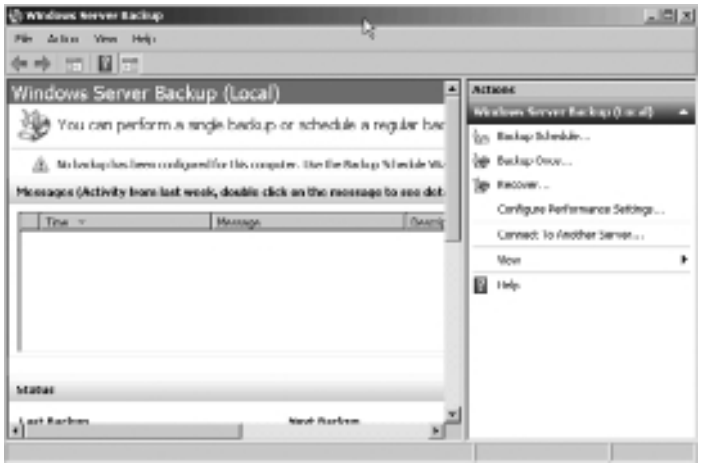

FIGURE 3.60
Windows Server Backup admin console

From the perspective of wanting to protect the data within the guest operating system using a simple interface for no additional cost (except for disk space), this is an acceptable solution. It is important to understand, however, that when the number of virtual machines begins to increase, the cost effectiveness of this solution begins to decrease. Unlike many of the commercial backup programs available today, the Windows Server Backup feature is geared more for backing up a small set of systems, local or remote. This lack of ability to group like systems into backup sets, see all backup clients in a single window, and have retention policies for backup media limits its use in large enterprise environments.

Many organizations choosing to adopt virtualization as the default deployment mechanism will already have a solid backup solution in place. If this is the case, the established solution may be suitable for the virtual environment as well. These solutions often require an agent to be loaded inside the guest operating system. Consult with the software vendor for any special requirements in a virtualized environment.

Virtual Machine Backup Solutions

As a complement to vSphere, VMware has introduced VMware Data Recovery (Figure 3.61). VMware DR creates backups of complete virtual machines, manages backup data based on retention policies, and uses deduplication technology to allow for the complete protection of a virtual machine with minimal storage consumption. VMware DR is an appliance-based solution and uses a vSphere client plug-in to manage and restore backups across your virtual environment. VMware Data Recovery provides the following features:

- vSphere integration using Open Virtualization Format (OVF) appliance and vSphere client plug-in
- Data deduplication automatically performed during every backup operation
- Volume Shadow Copy Service support for 32-bit and 64-bit Windows operating systems (Windows Server 2003, Vista, Server 2008)
- File-level restore within the VM with the optional File Level Restore (FLR) Client feature

FIGURE 3.61
VMware Data Recovery console

VMware Data Recovery provides a seamless mechanism for backing up the entire contents of the Windows virtual machines in your environment. The ability to support features such as VSS and data deduplication make it extremely attractive for Microsoft Windows–based virtual machines. With a full Windows 2008 install nearing 5 GB, the space savings can add up quickly. The following are a few points to keep in mind about VMware DR:

- VMware DR uses the vSphere licensing infrastructure, and valid licensing includes Essential Plus, Advanced, Enterprise, and Enterprise Plus.
- A single VMware DR appliance can protect a total of 100 virtual machines.
- A single VMware DR appliance can support two deduplication stores of up to 1 TB in size; this is a recommended maximum for optimal performance.
- Backups can be stored on any virtual disk supported by ESX, VMDK, RDM, and NAS, as well as CIFS; however, the preferred method is RDM or VMDK.
- Windows virtual machines must have VMware Tools installed and running for VMware DR to use the Volume Shadow Copy Service.

Third-party solutions also exist for complete virtual machine backup. Features to look for in a virtual machine backup solution include the following:

- VMware vStorage API integration
- Management console or vSphere Client integration
- Concurrent VM backups to decrease backup time
- Incremental backup
- Deduplication and/or compression
- VMware DRS/HA/VMotion compatible
- File-level restore

Summary

With the information and exercises provided in this chapter, you can begin building a solid foundation for your virtual Windows environment. You took a detailed look into the support considerations and guest OS requirements when virtualizing Windows Server 2008. The vendors you rely on to provide you with reliable software make great efforts to test different configurations using various hardware and software alongside their products. The outcome of these tests becomes part of the support matrices they provide. As mentioned previously, it is important to check up on these matrices every few months, especially any time you are implementing new (virtual) hardware or software. Once you understand the supported configurations and requirements, you can begin building those considerations into your baseline templates and processes. After following the exercises in this chapter, you have a golden Windows Server 2008 R2 template that will be a stepping stone for the chapters that follow.

As a vSphere administrator, you are familiar with the high-availability technologies that VMware provides such as HA, DRS, FT, and VMotion. As a Windows administrator, you understand the close ties that Microsoft Clustering Services and Failover Clustering have with your critical applications. Because of this relationship, there are times when you may need to deploy clustering within the guest operating systems. We covered the requirements and recommendations for deploying a failover cluster within a vSphere environment that is enterprise-ready.

As you make your way through the remaining chapters, you will be presented with best practices for deploying some of the most business-critical applications in use today while running in a VMware vSphere environment. The success of these deployments would not be possible without first establishing an efficient and consistent method for deploying the virtual machines that will run these applications.

Chapter 4

Virtualizing Active Directory

With the release of Windows 2000 Server, Microsoft established the Windows operating system as enterprise-worthy. There was a shift from the flat directory services that did not allow customization, could not scale to enterprise levels, and had inherent security issues. Active Directory Services resolves all of the shortcomings of the Windows NT operating system by allowing schema customizations, flexible design and deployment options, hierarchal framework, security, and better availability and recovery than its predecessor. With this new era in Microsoft directory services, many enterprises have deployed Active Directory Domain Services on physical hardware throughout the world.

As new versions of the Windows Server operating system come to the market, Active Directory is continuously being improved, and the hardware requirements are changing. This in part is providing enterprises with the opportunity to evaluate virtualization to reduce costs, power, cooling, and data-center space associated with deploying and refreshing hardware. As more and more system administrators look to virtualization to ease their daily operation woes of managing, provisioning, and updating operating systems, virtualization allows them to focus their efforts more on strategic planning and execution. System administrators can now get out of the daily grind of fighting fires. In addition, virtualizing Active Directory allows greater flexibility of testing and planning. You no longer have to order, provision, and configure physical servers to test migration or upgrade strategies.

Virtualization makes testing schema changes before you roll them into production environments easy and much more predictable. For example, Microsoft Exchange is one application that modifies the Active Directory and makes more than 100 changes/additions to the Active Directory schema to prepare for the Exchange messaging system. Although these schema updates and modifications have been thoroughly tested by Microsoft, few system administrators are willing to roll these changes right into production without investigating the impact these modifications might have on the production environment. Virtualizing Active Directory can help ease a system administrator's anxiety by facilitating a clone or snapshot of the production environment, which then can be used as a test and development environment for testing these schema modifications as well as operating system updates and for strategic disaster recovery and business continuity planning.

In this chapter, we will cover the following:

- The considerations for virtualizing Active Directory on the vSphere platform
- The importance of time synchronization and system-state backups
- The virtual machine build and configuration

We will also build a multinode proof-of-concept Active Directory forest in a lab configuration that will employ the topics discussed in this chapter.

Characterizing Your Active Directory Workload

In general, the Active Directory workload is an excellent candidate for virtualization and should not pose much of a challenge for most corporations. That said, you will need to be mindful of your current Active Directory deployment and monitor the replication traffic, server resource utilization, and additional services that are included as part of Active Directory. Windows Internet Name Service (WINS) if deployed, Domain Name Service (DNS), and global catalog (GC) servers for Exchange Server can place extra demands on server resources.

In addition, you must ensure that your current Active Directory forest is servicing clients in an accurate and timely manner. Failure to properly characterize your existing deployment can transition into the virtual design and deployment, thereby causing service delays and degraded performance in applications that might require authentication and directory lookups. Services such as DNS, Exchange, and SQL Server could experience performance degradation if Active Directory is not performing efficiently. Using such tools as ReplMon and RepAdmin can help you ensure that your domain controllers are replicating to their partners within the Active Directory forest. We will cover the use of ReplMon and RepAdmin later in the chapter.

Furthermore, monitoring the system event logs can provide indicators of degraded performance and service failures. These warning and error messages should be resolved in a timely manner. In addition, the use of the system performance monitor can give you insight into what resources are in use and where there might be deficiencies. You should monitor your domain controllers during normal and peak usage. Through daily, weekly, and monthly monitoring, you can ensure that your Active Directory domain controllers are healthy and properly servicing users and applications in a timely manner. Smaller directories tend to use fewer resources than larger directories that might contain the aforementioned services.

It is important that you provide the necessary resources such as virtual CPU, memory, and proper disk layout and configuration to the virtual deployment to ensure that your virtual deployment of Active Directory meets or exceeds the current business service-level agreements (SLAs) in place for your organization. By using the system performance monitor, you will be able to tell what resources are needed, if any, to ensure that your domain controllers are operating in an efficient manner. Failure to provide the required resources to the virtualized domain controllers will cause performance degradation with any authentication, DNS lookups, and password change requests and in some cases may cause domain controllers to become so far out of synchronization that they will no longer replicate with other Active Directory domain controllers. This can cause serious performance degradation as well as corruption within the Active Directory forest.

To recap, the following are some of the most important considerations when virtualizing your Active Directory forest:

- Characterizing your Active Directory workload
- Choosing between 32-bit or 64-bit virtual domain controllers
- Configuring the network
- Placing and configuring the vSphere host

- Synchronizing the Active Directory time
- Backing up and recovering Active Directory
- Securing your domain controller virtual machine files
- Monitoring replication

This chapter will cover these considerations for virtualizing Active Directory.

32-bit or 64-bit Active Directory Domain Controllers?

With Microsoft Windows Server 2003 through Windows Server 2008, you have the option of deploying your Active Directory virtual machines on either the x86 (32-bit) or x64 (64-bit) Windows Server platforms. With the introduction of Windows Server 2008 R2, the x86 platforms are no longer available. This shift in platform allows greater scalability of the operating system and the applications running on it. The 64-bit platforms allow for more memory allocation and also allow you to trade I/O through the use of larger memory caches. This enhanced caching ability results in more information being held in memory and a lesser need to access that information from the disk subsystem. This allows for a better overall operating system and application experience.

64-bit

If using the x64 version of Windows Server 2003 through Windows Server 2008 or R2, you will be able to contain the entire Active Directory database within the operating system's memory cache. vSphere affords a single virtual machine of up to 255 GB of memory. On 64-bit Windows, employing 32 GB of RAM to a virtual machine will accommodate a database of approximately 3 million users, not including Active Directory security or distribution groups, all in the operating system's memory. Caching the entire Active Directory database in the 64-bit Windows memory cache will avoid performance hits related to certain disk operations. For a virtual machine, that is a domain controller, the search, query, and certain update operations will generally benefit significantly from caching. Write operations such as delete, update, and modify will always incur a slight penalty because of the need to write changes to the disk subsystem. This is inherently true regardless of whether a domain controller is running on a physical or virtual machine.

32-bit

If deploying Active Directory on the Windows Server x 86 platforms, there will be limited benefit for filling the memory cache for customers with large directories; tests have shown that a database of about 2.75 GB can fit entirely into the operating system's memory cache on the Windows Server 2003 x86 (32-bit) editions. That will roughly translate to about 100,000 users or objects, not including Active Directory security groups, distribution groups, and other directory objects. If deploying large directories, you should consider deploying the Windows Server x64 platform editions for your Active Directory domain controller deployments. Larger enterprises that can't or choose not to deploy the 64-bit versions of the Windows Server 2008 platforms will be limited to the 4 GB operating system limitations and, if the entire directory cannot be cached, will need

to rely more on the shared disk subsystem. More domain controllers will need to be deployed in order to properly distribute the load to ensure that directory operations are not degraded.

Network Configuration

The network configuration of your vSphere hosts will play a large part in whether your virtualized Active Directory deployment is successful. Because the virtual machines will communicate through the physical network adapters, it is important that you configure your vSphere hosts with the proper number of network cards to ensure that you maintain the flexibility in your virtual to physical network design.

An example of this flexibility could be to use VLANs to separate your replication traffic from your normal virtual machine traffic. Providing a dedicated pipe for replication can guarantee as well as enhance the performance of Active Directory because when you provide dedicated bandwidth, you don't contend with normal network traffic and you ensure that network congestion will not produce replication delays or outages.

Selecting the appropriate virtual network adapter, physical network infrastructure, physical network bandwidth, and throughput will play a large part in the synchronization of the Active Directory forest. As we mentioned earlier, replication plays a large part in the synchronization of Active Directory.

Replication is how the directory is kept up-to-date as new objects such as printers, file shares, users, and groups are added to the directory. These objects as well as password changes from users must be replicated in a timely manner. For example, if you disable a user account and replication is not occurring because of an improper network configuration or physical network problems, that disabled account information may not be replicated in a timely manner, thereby causing a security breach in that the disabled user might be able to log in and access information that they are no longer supposed to be able to access. This can also occur for users who may be transitioning to a new job function.

For example, say John is a senior member of the human resources department and has access to employees' personal information. This information is stored on mapped drives that are automatically assigned when John logs onto his desktop. Then, John is promoted to senior marketing director and no longer requires access to this personal information. If Active Directory replication is experiencing problems because of physical network timeouts and replication is not succeeding, for a short window John could still have access to those mapped drives with the personal information on them because the changes made to disable John's access to this information was not replicated throughout the Active Directory. This can also cause a security breakdown for even active users.

You can use vCenter Maps view to analyze network infrastructure as it appears to your vSphere infrastructure deployment. These views allow you to verify that you have connectivity to the appropriate network interface controllers (NICs) and switches and will help you verify that your network paths are valid and functioning properly. Replication is a key component in maintaining a healthy Active Directory; proper selection of virtual network adapter, availability through network adapter teaming, separate virtual switches, and a properly configured physical network will ensure that Active Directory stays healthy. Consult VMware's best-practices guides on networking for information about developing a robust network infrastructure to support your virtual environment.

TIP For more information on vSphere v4.0 Best Practices, visit the VMware website at `www.vmware.com/resources/techresources/10041`.

vSphere Host Placement and Configuration

You need to consider where to place vSphere hosts based on your existing Active Directory configuration. In a physical Active Directory deployment domain, controllers should be located close to users, and this should also hold true for your virtual Active Directory deployment. Placing domain controllers close to users will facilitate local logons and prevent expensive wide-area network (WAN) bandwidth utilization by user logon traffic. Domain controllers should be located close to users. For smaller remote offices, domain controllers could be configured as part of a remote office/branch office (ROBO) solution that might include a file server and other application servers such as Exchange or SQL Server.

Virtualization can help simplify the build-out and deployment of remote/branch offices. These configurations can be designed and built out at the corporate office, and policies can be applied during the build-out. When the remote/branch office server is completed, the physical server hosting the ROBO solution can then be shipped to the remote/branch office, and less-technical personnel can connect the power and network to the machine.

Using vCenter, the remote/branch office solution can be managed from the corporate office. A subset of security privileges can be granted/delegated to the more technical personnel at the remote/branch office, allowing them to perform some of the more privileged functions at the ROBO such as resetting user passwords, printer management, and so on. Software and patch updates can be done remotely at the corporate office.

For the most part, domain controllers should be located where there might be more than 25 users. This will eliminate the need to authenticate over the WAN. To increase security, domain controllers in remote offices could be deployed as read-only domain controllers (RODCs). This type of domain controller only accepts changes from other writable domain controllers and provides a secure environment for remote office/branch office solutions.

When deploying Active Directory domain controllers, follow the Microsoft best practices, and deploy a vSphere host where you will deploy the virtualized Active Directory domain controller. Because Active Directory is a great candidate for virtualization, some older hardware may be able to be repurposed for your remote office/branch office deployments. When configuring your vSphere hosts to run Active Directory, ensure that you take into account the size of your directory and the Windows Server platform you will deploy for your domain controllers.

If you are virtualizing domain controllers that host a very large directory, take into consideration the size of your directory. You can usually determine how large your directory is by examining the attributes of the `NTDS.DIT` file in the `C:\Windows\system32` directory. This in addition to any other application CPU, memory, disk, and network requirements should give you a reference point for which to size the vSphere host.

Active Directory Time Synchronization

A large part of a successful Active Directory implementation is the proper planning of time services. In a virtualized environment, virtual machines that don't require CPU cycles don't get CPU cycles. For the typical Windows application in a virtual machine, this is not a problem. When virtualizing Microsoft's Active Directory, however, these idle cycles can cause significant time drift for domain controllers. Active Directory is critically dependent on accurate time-keeping, and one of the most important challenges you will need to address is how to prevent clock drift.

Kerberos v5

Microsoft uses Kerberos v5 as the authentication protocol, and time synchronization is critical to the proper functioning of this protocol. The implementation of Kerberos enhances the security in Microsoft Active Directory Domain services by keeping track of time-stamped issued tickets that are used for login authentication, server-to-server communications, and access to forest and domain-wide information and resources. These time-stamped authentication tickets generated by Kerberos are based on the workstation's time, and only a five-minute tolerance is allowed. If time drifts significantly, the authentication tickets will not be issued, will become outdated, or will simply expire, resulting in denied authentications or an inability to log into the domain to access network resources.

Controlling Time Drift

To control time drift and ensure that Active Directory time is synchronized throughout the Active Directory forest, you need to carefully consider the most effective way of providing accurate time to domain controllers. The domain controller with the PDC emulator role for the forest root domain is king and is the "master" time server for the forest. It's the root time server for synchronizing the clocks of the PDC emulators in each domain of the forest and, ultimately, all Windows computers in the forest. This PDC emulator can be configured to use an external source for its time. By modifying the defaults of this domain controller's role to synchronize with an external stratum-1 time source, you can ensure that all other DCs and workstations within the domain will be accurate. *Stratum* is used to refer to how close a computer is located to a time source. Stratum-0 devices are normally atomic or GPS clocks that are connected to computers. These computers then provide time services to stratum-1 devices that are normally computers or network devices. So, when you synchronize the PDC emulator to an external stratum-1 time source, you are in fact synchronizing to a computer on the Internet that is receiving its time services from an atomic or GPS clock. Figure 4.1 shows a representation of the Internet time hierarchy.

It is important to understand the relationship between the time source used by clients, member servers, and domain controllers. Figure 4.2 shows the source for the various entities. The PDC emulator role is an operations master role that advertises itself as the primary domain controller (PDC) to workstations, member servers, and domain controllers running earlier versions of Windows. Only one domain controller in each domain in the forest contains this flexible single master operations role.

The following sections cover some suggested ways to support this time-sensitive environment in the vSphere environment.

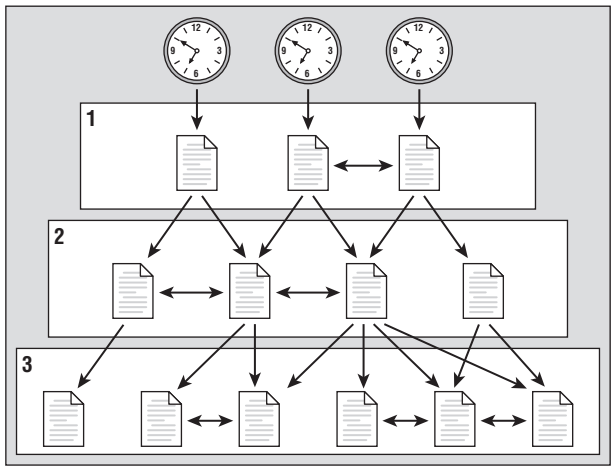

FIGURE 4.1
Hierarchy of Internet time sources stratum 0 (atomic or GPS clocks) at the root level

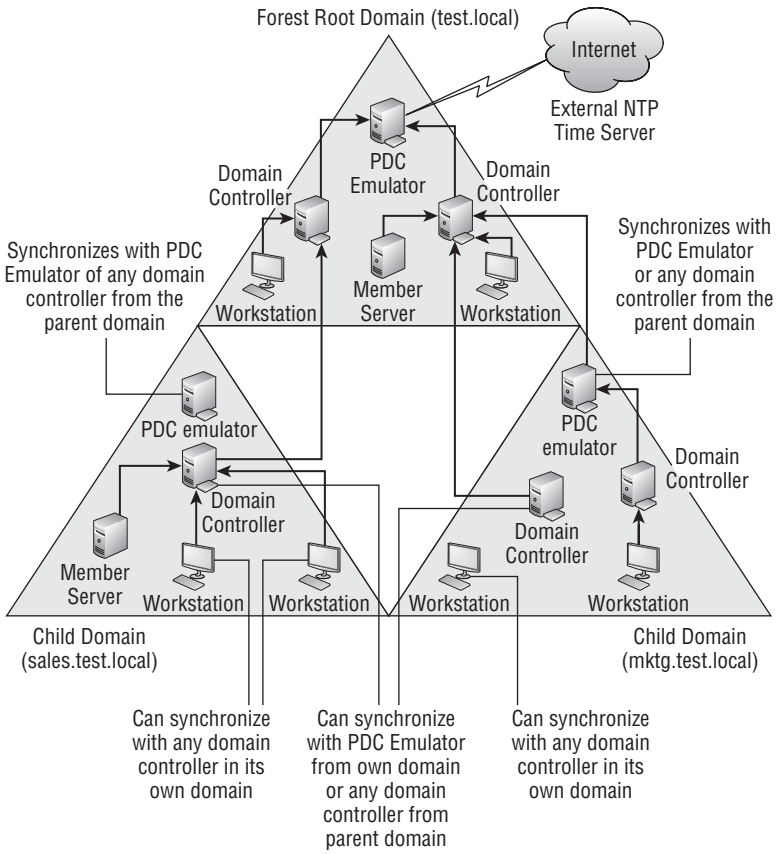

FIGURE 4.2
Synchronizing the forest root domain PDC emulator role to an external time source ensures all other DCs will be accurate.

Option 1: Using Group Policy and WMI Filtering to Modify the Windows Time Service

The first option is to use Group Policy to modify the Windows time service to synchronize with an external stratum-1 time source and link the policy to a WMI filter created to apply only to the PDC emulator FSMO role holder. By querying for the PDC emulator role and applying the Group Policy to that role, you can always ensure that the PDC emulator role will synchronize with an external stratum-1 time source and provide accurate time to all members of the forest/domains. We will show how to do this in the "Active Directory Proof of Concept" section later in the chapter, so we will not go into great detail about this option here.

Option 2: Using the Windows Time Service for Synchronization

The second option is to use the Windows Time Service and not VMware Tools synchronization for the forest root PDC emulator. This requires configuring the forest PDC emulator to use an external time source. The procedure for defining an alternative external time source for this "master" time server is as follows:

1. Modify registry settings on the PDC emulator for the forest root domain.

   a. In the key HKLM\System\CurrentControlSet\Services\W32Time\Parameters\Type, change the Type REG_SZ value from NT5DS to NTP.

      This determines from which peers w32time will accept synchronization. When the REG_SZ value is changed from NT5DS to NTP, the PDC emulator synchronizes from the list of reliable time servers specified in the NtpServer registry key.

   b. In the HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Parameters\NtpServer key, change the NtpServer value from time.Windows.com,0x1 to an external stratum-1 time source—for example, timekeeper.isi.edu,0x1.

      This entry specifies a space-delimited list of stratum-1 time servers from which the local computer can obtain reliable time stamps. The list can use either fully qualified domain names or IP addresses. (If DNS names are used, you must append ,0x1 to the end of each DNS name.)

   c. In the HKLM\System\CurrentControlSet\Services\W32Time\ key, change AnnounceFlags REG_DWORD from 10 to 5.

      This entry controls whether the local computer is marked as a reliable time server (which is possible only if the previous registry entry is set to NTP as described earlier). Change the REG_DWORD value from 10 to 5 here.

2. Stop and restart the time service:

```
net stop w32time && net start w32time
```

Microsoft Knowledge Base article 816042 provides detailed instructions for this process.

WARNING Choose one of the time synchronization options presented. The two or more time synchronization methods are not aware of the other's existence, and attempting to employ multiple options at once can cause erratic behavior or time flip-flopping.

Option 3: Modifying the Windows Time Service

The third option is to modify Windows Time Service through the use of a domain controller's Group Policy for all DCs, implement the NTP daemon from within the Service Console pointed to an external stratum-1 NTP source, and use VMware Tools time synchronization for the domain controller's virtual machines. Here is the procedure:

1. Implement a domain controller's Group Policy to modify the registry as shown in Figure 4.3.

In this key change NTDS5 to NTP:

```
HKLM\System\CurrentControlSet\Services\W32Time\Parameters\NTP
```

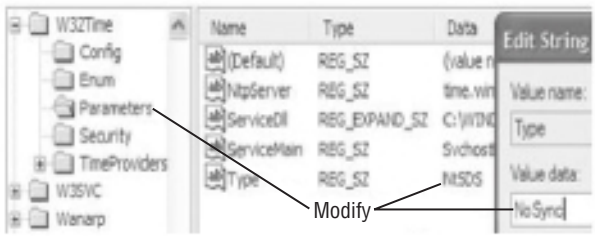

Figure 4.3 Modifying the registry for no synchronization

2. Enable the ESX server NTP daemon to sync with external stratum-1 NTP source as detailed in VMware Knowledge Base article 1339.
3. Use VMware Tools time synchronization (Options | Time synchronization between the virtual machine and the console operating system) within the virtual machine.

WARNING VMware Tools time synchronization is designed to play "catch up," not slow down. What this means is if your system clock is behind, VMWare Tools will catch up your system clock to the correct time, but if your system clock is fast, VMware Tools will not slow down the system clock.

Active Directory Backup and Recovery

You should always endeavor to perform server and system-state backups on a daily basis. If you cannot perform backups with this regularity, you should perform at least a weekly backup of the server and system-state data for your Active Directory domain controllers.

We recommend that you never restore system-state data older than 60 days because this is coming pretty close to the 60 day default tombstone expiration date for Active Directory objects within the forest. Your domain controller's system-state backup should consist of the SysVol information, boot and system files, the registry, and the Active Directory service.

When you use NTBackup, Windows Server Backup, or most third-party vendor programs, selecting to include the system-state option will ensure that you have all the necessary information to restore an Active Directory domain controller to a operational state in the event of failure. Never use differencing or snapshot disks when virtualizing Active Directory. This makes it too easy to improperly back up and restore Active Directory. Just simply performing a snapshot operation on a virtualized Active Directory domain controller can cause the domain controller to become corrupted.

Restoring a virtualized domain controller from a snapshot or an improperly backed up domain controller will not activate the internal Active Directory database checkpoint mechanisms that alert the Active Directory domain that a failed domain controller has been restored and placed back into operation. These mechanisms ensure that the restored domain controller can receive changes from other domain controllers in the forest/domain and replicate changes to other domain controllers in the forest/domain.

A domain controller that is improperly backed up via snapshots can cause corruption to the virtualized domain controller as well as corrupt the entire Active Directory forest if improperly restored. This data corruption might not be noticeable at first but could later cause catastrophic failure of the entire Active Directory forest, thereby causing you to have to rebuild your entire Active Directory forest. Performing regular, consistent system and system-state backups ensures the integrity of the Active Directory database by committing transactions and updating database IDs.

WARNING Never snapshot an Active Directory domain controller because doing so could cause corruption on that domain controller, which will spread to other domain controllers in the Active Directory forest causing catastrophic failure. Although Windows Server 2008 R2 does a great job at preventing corruption to domain controllers in the event of improper restore procedures, we recommend that you follow the Microsoft best practices for backing up and restoring Active Directory domain controllers.

During the system-state backups, these database IDs (invocation IDs) and update sequence numbers (USNs) are properly set and identified. Figure 4.4 shows the effects of improperly restoring a domain controller from a differencing disk or snapshot.

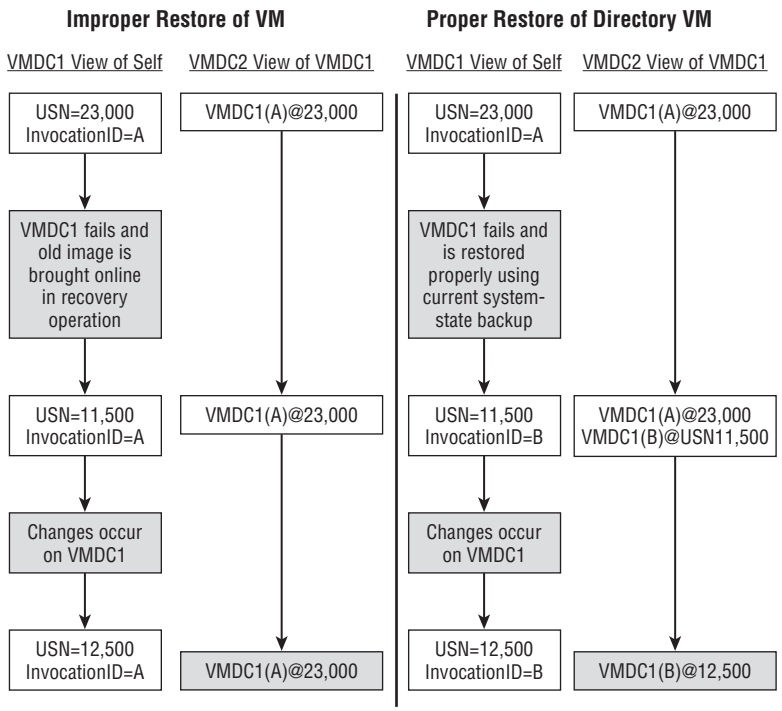

FIGURE 4.4 Improperly restored domain controller

In Figure 4.4 on the left, restoring a virtual machine from an old copy of a virtual disk does not update the database IDs or the USNs. As you can see, the databases invocation IDs are not synchronized. DC2 will never accept inbound replication of the changes since it already has knowledge of the objects that reside in VDC1's Active Directory database. Domain controllers are interested only in changes that are newer than the database they contain. Using the proper method, as shown on the right, ensures that missing transactions are restored correctly to the failed database once it synchronizes with its replication partners. This ensures that Active Directory replication is properly restored to the domain controller that was recently placed back into operation and that there is no corruption of the domain controller or the Active Directory forest. Using snapshots or disk imaging of any domain controller is not supported by Microsoft. To have a supported domain controller virtual machine, system state data backups must be performed to protect the domain controller and the Active Directory database.

WARNING Never restore an Active Directory domain controller from a snapshot or old virtual disk because this will introduce objects that may have already been deleted or tombstoned back into the directory and may cause corruption throughout the Active Directory forest/domain.

Developing well-defined and executed processes and procedures for disaster recovery is another important facet of virtualizing Active Directory. It is important to create and execute procedures for backing up a domain controller's system state data with regularity. You can use Window Server's NTBackup or third-party solutions that execute the appropriate Windows system backup APIs. For Windows Server 2008 and R2, there is the new Windows Server Backup program that can be scheduled to automatically back up the domain controller on a regular basis. There is a new option to perform a full or differential system-state backup.

Also new for Windows Server 2008 R2 is the ability to mount the system-state backup and verify as well as selectively restore information to the domain/forest. The backup solution that you choose should ensure it is capable of doing system-state backups. Remember, snapshots can potentially corrupt the Active Directory database without even committing or reverting to the snapshot. (See `http://support.microsoft.com/kb/888794`.)

Securing Your Domain Controller Virtual Machine Files

Providing secured access to domain controller virtual machines and their virtual files and disks can be leveraged through the use of VMware vCenter access roles. With the use of access controls in ESX 4/vCenter 4, you can give specified users and/or groups the level of privilege they require to maintain the domain controller virtual machines without affecting anything else within the vSphere environment.

The new, customizable roles give the administrator a great deal more flexibility than previous versions of VMware Virtual Center. In particular, the role-based access control (RBAC) features of the new vCenter provide greater control over virtual disk (.vmdk) file access. You can create role-based definitions that allow only the most trusted individuals in your organization access to the critical Active Directory data that runs your small to enterprise business operations.

TIP Use vSphere's RBAC features to enhance the security of your virtualized Active Directory domain controller files and virtual disks.

For instance, you could create a folder called vDomain_Controllers in the vCenter application and add the domain controllers into this folder, and under the Administrator ➢ Roles container you could create a role called Active Directory Administrators role, apply it to this newly created folder, and add only the trusted system administrator members for your organization.

These added users would be selected from the Active Directory user and group information at the Active Directory level. This would ensure that only your trusted domain administrators would have direct virtual machine (VM) access.

For example, say a new system administrator is hired and is also new to virtualization. By giving this new administrator view-only rights on the folder containing the virtualized domain controllers, you enable the administrator to get better acquainted with virtualization.

At the same time, you protect your virtualized domain controller files against accidental corruption and deletion. By giving trusted personnel the permissions to perform only their required job tasks, you can reduce the exposure to malicious and accidental misuse of authority in the virtual environment.

This also helps restrict the number of members of your organization who can gain access to information that they do not and would not normally require to perform their required job functions. Securing your virtualized domain controllers from unauthorized access will prevent accidental and malicious actions and save you major headaches when trying to recover your domain controllers as the result of these actions.

Monitoring Active Directory Replication

Accurate timekeeping is also essential for the replication process in a multimaster directory environment. Active Directory replication uses USNs to determine what changes to pull from its replication partners.

Domain controllers only look for records newer than their own. If simultaneous changes made on different domain controllers result in identical USN numbers, time stamps are used as one form of tie-breaker. In a multidomain controller or multidomain environment, the successful exchange of Active Directory database information with replication partners is vital.

Utilizing tools to determine paths of connectivity and site replication will ensure consistent and accurate databases. You must regularly monitor Active Directory replication to ensure those domain controllers are kept in synchronization within the Active Directory forest.

Failure to monitor Active Directory on a consistent regular basis will result in missed opportunities to correct any replication errors that might go unnoticed as a result of not monitoring. These errors could lead to catastrophic failure of the entire forest. This can happen regardless of whether you maintain physical or virtualized domain controllers. Active Directory ReplMon (`replmon.exe`) can provide very detailed information about the state of affairs within your Active Directory infrastructure.

This tool can show you a great deal of information about connections, updates, replication status, FSMO roles, and so forth. Additionally, it can graphically show the relationship between domain controllers. This tool indicates when there is a replication problem, as illustrated in Figure 4.5.

Figure 4.6 shows the output of the RepAdmin tool, which we will cover in more detail in the "Active Directory Proof of Concept" section of this chapter. From the output you can determine that Active Directory replication is functioning properly.

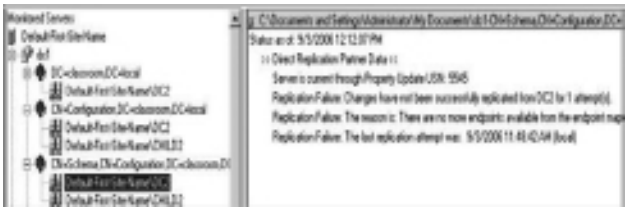

FIGURE 4.5
Detecting Active Directory replication issues

FIGURE 4.6
Monitoring Active Directory replication

Tools

Microsoft provides numerous tools to monitor the health of Active Directory. For users deploying Active Directory on Windows Server 2008 and R2, the following are some tools you can use to test and verify the state and health of the Active Directory and its replication:

- Replication Monitor (ReplMon)
- Replication Administration (RepAdmin)
- Domain Controller Diagnostics (DCDiag)
- Domain Secure Channel (NlTest)
- Microsoft IT Environmental Health Scanner
- Active Directory Best Practices Analyzer (ADBPA; Windows Server 2008 R2 only)

Because we will talk about ReplMon and RepAdmin in some detail in the "Active Directory Proof of Concept " section, we will not discuss these two tools here. The Microsoft IT Environmental Health Scanner is available on the Microsoft website under the Downloads section. The DCDiag and NlTest tools are now included with the operating system and get installed when you install the OS.

DCDiag

The Domain Controller Diagnostics tool can give you detailed information about how your Active Directory domain controllers are performing (Figure 4.7). It can provide information such as the health of your Active Directory domain partitions, including critical information such as errors, warnings, and informational messages that are logged in the Windows system event log files. This information can provide you with detailed insight into Active Directory operational status as well as alert you to possible issues that may be occurring within Active Directory before they become critical issues. Other operations that are critical to the health of Active Directory are also reported such as file replication system status, FSMO role status, replication activity status, and the knowledge consistency checker (KCC) status. You should monitor this information on a regular basis to ensure that Active Directory is functioning properly.

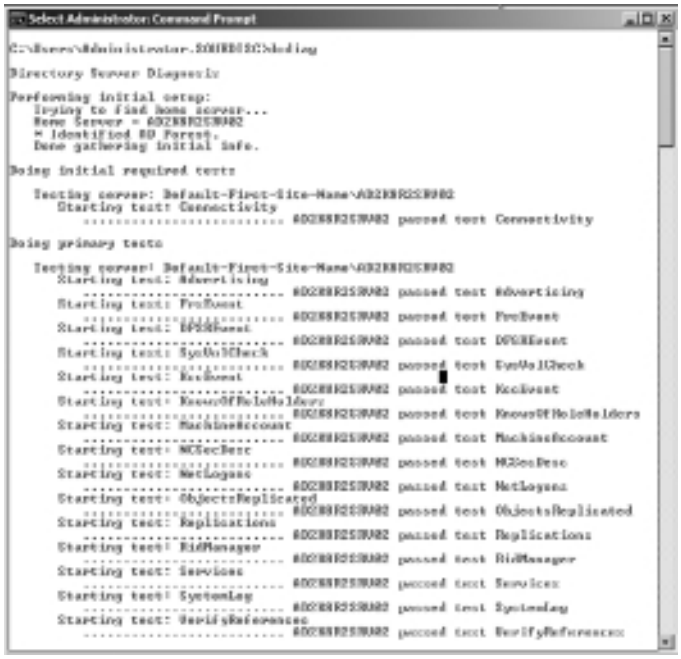

FIGURE 4.7
DCDiag partial display of Active Directory status

Figure 4.7 is the first part of the DCDiag output displaying the various critical operations within Active Directory along with their status. Any failures should be investigated immediately to ensure Active Directory remains operational.

In Figure 4.8, the second portion of the figure displays the Active Directory partition tests' status. This section will alert you if replication is not occurring or has failed or if there is corruption within the partitions displayed. As with the previous figure, any failures should be investigated immediately to ensure Active Directory health and functionality.

FIGURE 4.8
DCDiag completed display of Active Directory status

NLTEST

The Domain Secure Channel tool can display such information as domain controllers in the forest/domain, PDC emulator role holder, and logons to the domain controller. This tool can also allow you reset or change the secure channel password that is used for communication between domain controllers. This tool can alert you when a domain controller is no longer communicating with its partner within the Active Directory. This tools also enables you to locate domain controllers that have the flexible single master operations role. This tool has many parameters and options for obtaining information about the operation of domain controllers within the Active Directory forest/domain. For more information on the NlTest diagnostics tool, issue the following command at a command prompt:

```
Nltest /?
```

This will list all of the parameters and their options. Figure 4.9 displays the output of some of the more common commands used with NlTest.

Administrator: Command Prompt
C:\>nltest /dsgetdc:
Flags: 0
Connection Status = 0 0x0 NERR_Success
There was no failure in the last update for all DC-specific DNS records
The command completed successfully
C:\>nltest /sc_query
Number of attempted logons: 0
The command completed successfully
C:\>nltest /dsgetdc:sourdisc.com
DC: \\
Address: \\192.168.100.2
Dom Guid:
Dom Name: sourdisc.com
Forest Name: sourdisc.com
Dc Site Name: Default-First-Site-Name
Our Site Name: Default-First-Site-Name
Flags: GC DS LDAP KDC TIMESERV WRITABLE DNS_DC DNS_DOMAIN DNS_FOREST CLO
SE SITE FULL SECRET
The command completed successfully
C:\>nltest /dclist:sourdisc.com
Get list of DCs in domain 'sourdisc.com' from '\\
[DS] Site: Default-First-Site-Name
[PDC] [DS] Site: Default-First-Site-Name
[DS] Site: Default-First-Site-Name
[DS] Site: Default-First-Site-Name
The command completed successfully
C:\>nltest /domain_trusts
List of domain trusts:
0: SOURDISC sourdisc.com (NT 5) (Forest Tree Root) (Primary Domain) (Native)
The command completed successfully
C:\>
C:\>nltest /sc_change_pwd

FIGURE 4.9 Output from the NlTest diagnostics tool

What to Monitor

There are five Active Directory partitions that must be replicated fully and in a timely manner to ensure your Active Directory is healthy. There are two naming contexts that replicate throughout the forest: the schema and the configuration naming contexts. There is one of each of these naming contexts for the Active Directory forest. The other two naming contexts are domain-specific, and there could be many of these depending on the number of domains you have within your Active Directory forest. These naming contexts are the Domain and DomainDNSZones partitions.

Figure 4.6 (shown earlier in the chapter) shows you the RepAdmin utility. You will notice that replication is occurring normally for the Active Directory shown in the figure. The replication has completed successfully for the five Active Directory partitions.

NOTE The ReplMon, DCDiag, and NlTest tools are also available for Windows Server 2003 if you are installing Windows Server 2003 Active Directory domain controllers. These tools can be installed from the Support Tools directory of the Windows Server 2003 CD.

Active Directory Proof of Concept

After you have determined that your Active Directory workload is an excellent candidate for virtualization and have observed that it exhibits no abnormal behavior, you can proceed to a proof of concept. In this proof-of-concept scenario, we will show how to create the Active Directory virtual machine, synchronize time on the PDC emulator with the external stratum-1 time source, back up the Active Directory virtual machine, and verify Active Directory replication. Let's get started.

Creating and Configuring the Active Directory Virtual Machine

Based on the observed Active Directory workload in the physical environment, the domain controllers will have the following virtual hardware for the virtual machines:

- A virtual machine with two virtual CPUs (vCPUs)
- 2 GB of memory
- One VMXNET3 network adapter
- Three virtual disks (OS, Active Directory database, and logs)
- Two virtual SAS adapters (one for the OS and logs and one for database)
- Microsoft Windows Server 2008 R2

Figure 4.10 shows some starting points for sizing Active Directory domain controllers. In this figure, we show the virtual CPU (vCPU), memory, and disk requirements for small and larger directories. When sizing Active Directory, start with these suggestions, and increase vCPU, memory, and disk resources as required by your Active Directory domain controller workloads.

NOTE To ensure performance of the Active Directory VM, monitor the number of virtual machines that will be placed on the data stores containing the Active Directory domain controllers because this will ensure that there is no contention for disk resources at the storage array and within the virtual machines.

FIGURE 4.10
Suggested starting points for sizing Active Directory Domain controllers

Active Directory Domain Controller	
Hardware (Virtual)	VMware EXS 4.0 (vSphere)
Processor	2 vCPU (smaller directories start with a single vCPU) 2 vCPU (larger directories start with 2vCPU and increase as required)
Memory	4 GB (smaller directories start with ~1–2 GB) ***100,000 users can require up to 2.75 GB of memory to cache directory (x86)*** ~16 GB or greater (larger directories start with 16 GB) ***3 Million users can require up to 32 GB of memory to cache entire directory (x64)***
Network Adapter	Enhanced VMXNet3 Physical NICs should be teamed and plugged into separate physical switches in your network infrastructure
Storage Configuration	
vSCSI Controller 0 vSCSI Controller 1	Disk1 (C:) OS ~16 GB or greater (depending on OS version) Disk2 (D:) Database ~16 GB or greater for larger directories Disk3 (L:) Log files ~4 GB or 25% of the database LUN size

Utilizing separate virtual disks and vSAS controllers for the OS, logs, and database will allow you to separate the different I/O required by the database and the logs. Also, when using multiple vSCSI or vSAS adapters, the VMkernel will provision separate data structures for each adapter, giving you the ability to enhance the I/O throughput by allowing you to put more data "in

flight." Figure 4.11 shows that we placed our disks in independent persistent mode to prevent the creation of snapshots of the virtual disks.

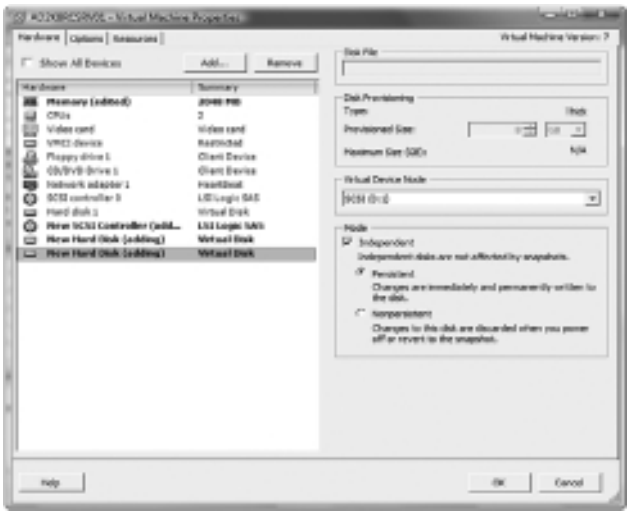

FIGURE 4.11
Active Directory virtual machine configuration

INSTALL WINDOWS SERVER 2008 R2

After configuring the virtual machine, you power on and proceed to install the Windows Server 2008 R2 operation system. On the initial boot screen, select the edition of the operating system that you want to install (Figure 4.12). Although we could have installed the Windows Server 2008 R2 standard edition, for this example we chose to install the 64-bit enterprise version at the request of Source Discovery.

NOTE In Chapter 3 we detail how to create and configure virtual machines and how to create templates of those created virtual machines. We recommend you use your created template to provision your domain controllers because this will save significant time in the provisioning process.

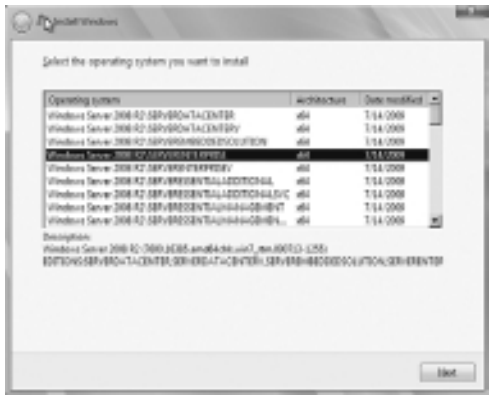

FIGURE 4.12
Selecting the operating system version to install for Active Directory

At this point, accept the end user license agreement (EULA), and continue with the installation of the Windows operating system. Next, you come to the installation type selection, where you select the Custom option to install a new clean version of the operating system onto the virtual machine. When we provisioned storage for our virtual machine, we allocated the storage for Active Directory services as shown in Table 4.1.

Next you come to the disk selection screen of the installation wizard (Figure 4.13); for this example, we're choosing to install the OS onto Disk 0. This will hold all of the files necessary to boot and run the operating system and Active Directory services.

TABLE 4.1: Active Directory Domain Controller Disk Sizes

DISK	SIZE
Operating system	32 GB
Active Directory database	32 GB
Active Directory transaction logs	8 GB

FIGURE 4.13
Selecting the disk for operating system installation

The two remaining disks will be used to host the Active Directory services database and log files. It is a Microsoft best practice to separate the Active Directory databases and log files onto separate disks. For more information on Active Directory disk layout, see the Microsoft website. After you make your disk selections, the installation wizard proceeds to install the Windows operating system onto the selected disk.

After the installation has completed and the virtual machine has rebooted, you must log into the operating system. Upon your first login, you will be prompted to change the newly installed server password (Figure 4.14). Depending on your domain password policy, you may have to specify a password that contains letters, numbers, and other characters. This is known as a *strong password* requirement. You can relax this requirement by modifying the domain policy, but we advise against it because it is a Microsoft best practice to utilize strong passwords to ensure the security of your domain controllers.

FIGURE 4.14
Changing the password for the initial login to the Windows server

INSTALL VMWARE TOOLS

After logging into the server, the first action you need to take is to install VMware Tools onto the server. Installing VMware Tools will greatly enhance the video, keyboard, and mouse functions on the virtual machine as well as install the VMXNET3 network adapter for use. You get to this by selecting VM ➢ Guest ➢ Install/Upgrade VMware Tools (Figure 4.15).

WARNING If you do not install the tools, you may find accessing the virtual machine frustrating. When using the virtual machine's console, video response may be jerky and in some cases unresponsive. Mouse movements may be unpredictable, and your keyboard response might have you typing multiple characters even though you may have typed only one. Installing the VMware Tools is a must for any virtual machines that are hosted on any VMware product.

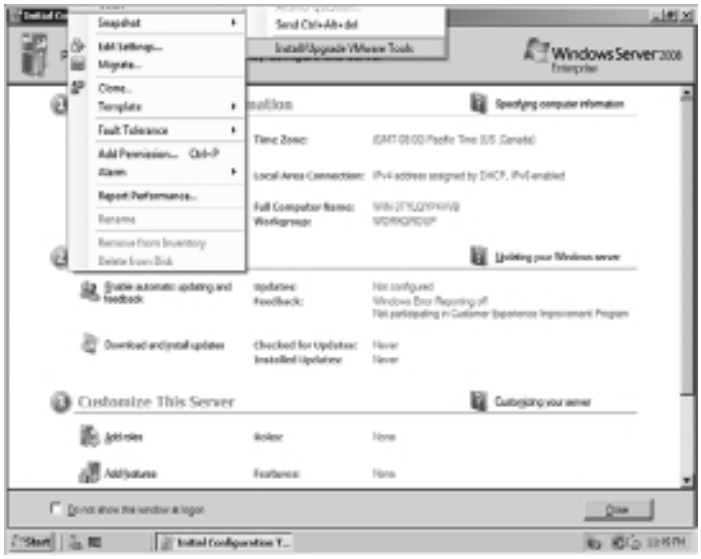

FIGURE 4.15
Installing VMware Tools on the server

After the installation of VMware Tools is completed, you are prompted to reboot. Reboot the server, and upon reboot, you should be able to configure the vmxnet-3 network adapter.

Configure the Network

When configuring the server IP address, it is recommended that you use static addresses for your domain controllers. As a best practice, in the physical as well as the virtual environment, servers and especially domain controllers should not be configured using the Dynamic Host Configuration Protocol (DHCP). If there is a failure with the DHCP server in the domain, your domain controllers will not obtain a TCP/IP address and therefore will not be able to communicate on the network. We will not bore you with the network setup because you should be familiar with the configuration of the TCP/IP settings for your servers. Figure 4.16 shows you the network properties that were configured for our virtual machine.

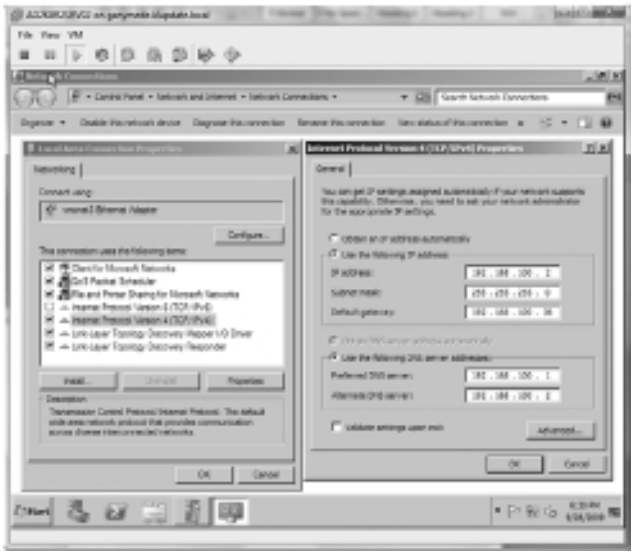

Figure 4.16 Configuring the TCP/IPv4 properties of a virtual machine

Source Discovery does not utilize TCP/IPv6, so we disabled the IPv6 protocol for the vmxnet-3 adapter and configure the IPv4 settings for the virtual machine. After all the settings are entered, click OK to activate the entered settings for this adapter. You can then proceed to enter the computer name for the virtual machine and reboot the virtual machine.

TIP Some folks prefer to enter the name and join the domain at the same time. We have experienced from time to time—on physical servers as well virtual machines—a situation where we entered the machine name and joined the domain and then when the server was rebooted, we couldn't log into the domain because the trust account was not properly created. This is why we suggest you enter the name of the server and reboot and then join the domain and reboot again. It's a pain, but it works every time.

Installing the Active Directory Service

After you get the machine into the domain, it is time to install the required services to make this virtual machine a domain controller in the Source Discovery (sourdisc.com) domain. When preparing the virtual machine for the installation of Active Directory using the Server Manager's Add Roles Wizard, we installed the .NET Framework v3.51 and the Active Directory Services role onto the virtual machine. For the Active Directory Domain Services role, we selected the Active Directory Domain Services role and proceed with the installation on this virtual machine (Figure 4.17).

TIP View the Before You Begin screen to ensure you meet the prerequisites for the roles that you are attempting to install.

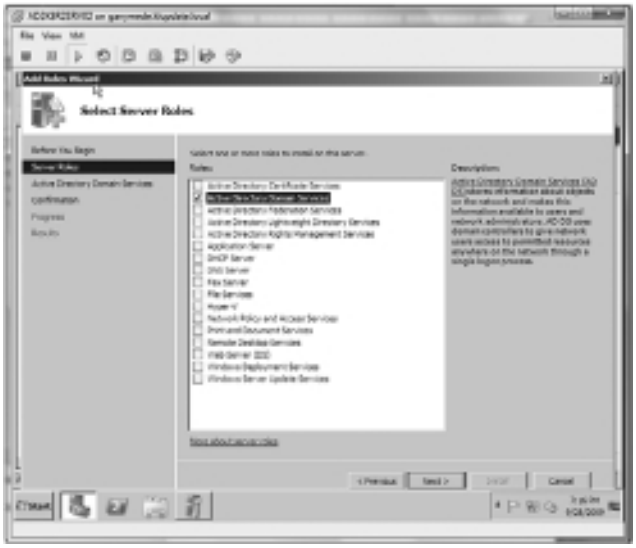

FIGURE 4.17 Installing the Active Directory Domain Services role

After the roles has successfully installed on the virtual machine, you can click the link in the Installation Results window to start the installation of Active Directory Domain Services on the virtual machine (Figure 4.18); or click Start ➢ Run, type **`dcpromo`** in the Run box, and click OK to make this virtual machine an Active Directory domain controller.

During this wizard you determine whether you're joining an existing forest and just adding domain controllers or whether you're creating an entirely new forest. Because Source Discovery wants to preserve its existing forest, you opt to join the existing forest and add domain controllers to the sourdisc.com forest. The additional domain controllers will be installed on different hosts of the VMware DRS/HA cluster. Affinity using the separate VMs will be created and applied to the domain controllers so that all of the virtual domain controllers do not end up on the same host within the VMware DRS/HA. In this scenario, you show the technical staff how to create the affinity rule after all of the virtual domain controllers are installed and configured, and then you can move on to the configuration of Active Directory on this virtual machine.

Active Directory Domain Services Installation Wizard
Welcome to the Active Directory Domain Services Installation Wizard
This wizard helps you install Active Directory Domain Services (AD DS) on this server, making the server an Active Directory domain controller. To continue, click Next.
Use advanced mode installation
Learn more about the additional options that are available in advanced mode installation.
More about Active Directory Domain Services
< Back
Next >
Cancel

FIGURE 4.18
Starting the installation of the Active Directory onto the virtual machine

In Figure 4.19, we are opting to add a domain controller to the existing sourdisc.com domain. This will add redundancy and load balancing to the existing Active Directory forest while allowing us to realize the final end-state Active Directory virtualization goal. We will allow the additional virtualized domain controllers to run alongside the physical domain controllers in production for the determined amount of time before transferring the FSMO roles to the virtualized domain controllers. Once the roles have been transferred to the virtual domain controllers, the physical domain controllers will decommissioned, and the hardware (which is currently Core2Duo technology) will be repurposed for possible use in the VMware DRS/HA clusters as additional hosts.

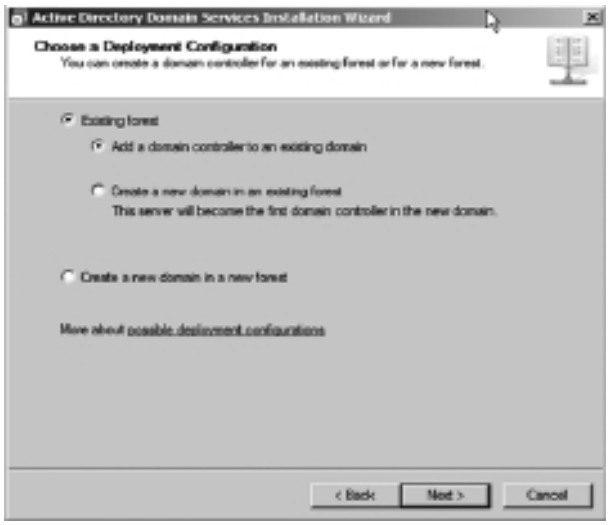
Active Directory Domain Services Installation Wizard
Choose a Deployment Configuration
You can create a domain controller for an existing forest or for a new forest.
Existing forest
Add a domain controller to an existing domain
Create a new domain in an existing forest
This server will become the first domain controller in the new domain.
Create a new domain in a new forest
More about possible deployment configurations
< Back
Next >
Cancel

FIGURE 4.19
Adding a domain controller to an existing domain

In Figure 4.20, we are selecting the sourdisc.com domain and choosing to provide alternate credentials. Because the virtual machine is not a member of the sourdisc.com domain, we will have to provide an account that has permissions to join the domain and install the Active Directory Domain Services onto the virtual machine. If the machine was already part of the sourdisc.com domain, you would only be required to have permissions to install Active Directory Domain Services onto the virtual machine.

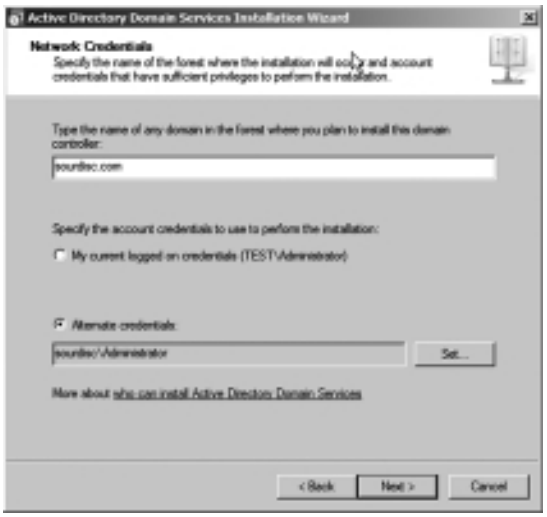
Active Directory Domain Services Installation Wizard
Network Credentials
Specify the name of the forest where the installation will occur and account credentials that have sufficient privileges to perform the installation.
Type the name of any domain in the forest where you plan to install this domain controller:
sourdisc.com
Specify the account credentials to use to perform the installation:
My current logged on credentials (TEST\Administrator)
Alternate credentials:
sourdisc\Administrator
Set...
More about who can install Active Directory Domain Services
< Back
Next >
Cancel

FIGURE 4.20
Providing the domain name and credentials of domain to join

Source Discovery has opted to make all of the domain controllers at the corporate and remote data center global catalog servers for optimal performance for all of those applications that rely on the global catalog servers. They will also create one writable domain controller and one RODC at each remote office that will host an Exchange Server 2007 server. The primary reason for creating one writable and one read-only is that Exchange won't talk to a read-only domain controller or read-only global catalog server. For these sites, there will be one writable domain controller that will also be configured as a global catalog server (Figure 4.21).

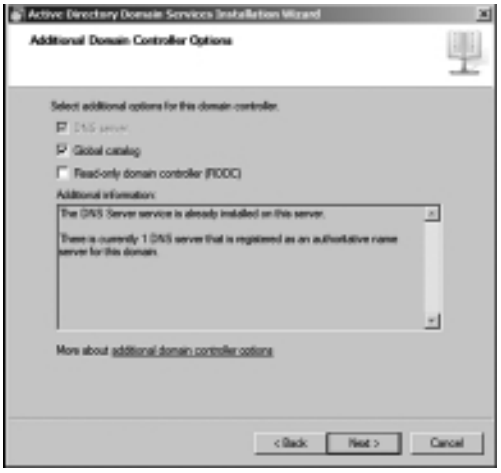
Active Directory Domain Services Installation Wizard
Additional Domain Controller Options
Select additional options for this domain controller.
DNS server
Global catalog
Read-only domain controller (RODC)
Additional information:
The DNS Server service is already installed on this server.
There is currently 1 DNS server that is registered as an authoritative name server for this domain.
More about additional domain controller options
< Back
Next >
Cancel

FIGURE 4.21
Selecting to install the global catalog onto the domain controller

In the presentation to the Source Discovery technical staff, it was communicated that the placement of the Active Directory database and log files play an important factor in the performance of Active Directory (Figure 4.22). Here we opted to place the Active Directory database and log files onto separate virtual disks that were provisioned out of different storage spindles/disks to provide maximum performance. Also, these data stores and virtual disks are thick-provisioned, meaning that all space is allocated at the time of creation of the data store or virtual disks. Microsoft has a best-practice recommendation of using thick provisioned nondifferencing virtual disks for the Active Directory database and log files. Best practice was implemented by using thick-provisioned disks for the Active Directory database and log files and by also placing them on separate virtual disks that are provisioned out of separate spindles at the storage array level.

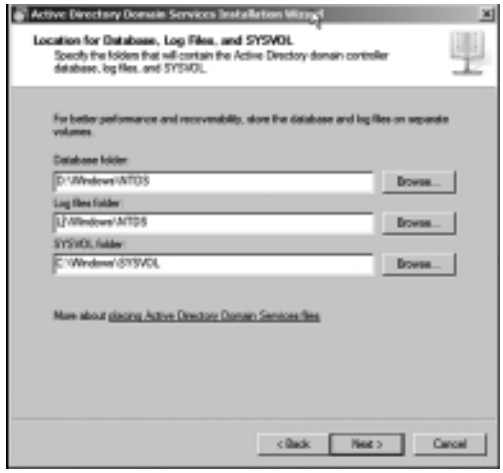
Active Directory Domain Services Installation Wizard
Location for Database, Log Files, and SYSVOL
Specify the folders that will contain the Active Directory domain controller database, log files, and SYSVOL.
For better performance and recoverability, store the database and log files on separate volumes.
Database folder:
D:\Windows\NTDS
Browse...
Log files folder:
L:\Windows\NTDS
Browse...
SYSVOL folder:
C:\Windows\SYSVOL
Browse...
More about placing Active Directory Domain Services files
< Back
Next >
Cancel

FIGURE 4.22
Selecting the location for the files critical to Active Directory

Here you can observe the various Active Directory checks and install operations. These will ensure that the Active Directory domain controller is installed and configured correctly (Figure 4.23), including checking for the presence of DNS; securing the security accounts manager (SAM); creating, configuring, and securing the database and log files; and creating, securing and sharing the SysVol folder used for replicating the domain public files and critical Active Directory information such as scripts and Group Policy. After installation of Active Directory Domain Services onto the virtual machine, you select to reboot the server. Upon reboot, you will log into the domain controller and proceed to configure the domain controller for best practices through the use of the Active Directory Best Practices Analyzer (ADBPA). This will ensure that your virtualized Active Directory domain controller is compliant with the Microsoft recommended best practices for deploying Active Directory.

Active Directory Domain Services Installation Wizard
The wizard is configuring Active Directory Domain Services. This process can take from a few minutes to several hours, depending on your environment and the options that you selected.
Configuring the local computer to host Active Directory Domain Services
Cancel
Reboot on completion

FIGURE 4.23
Active Directory installation and configuration of the virtual machine

Figure 4.24 shows the login screen for the newly created domain controller. Notice the domain user on the login screen. Here you will log in to the domain controller using the Administrator account for the sourdisc.com domain. Once you are logged into the domain, you use the Server Manager to run the Active Directory Best Practices Analyzer, as shown in Figure 4.25.

FIGURE 4.24
Domain Controller logon screen

FIGURE 4.25
Active Directory Best Practices Analyzer results

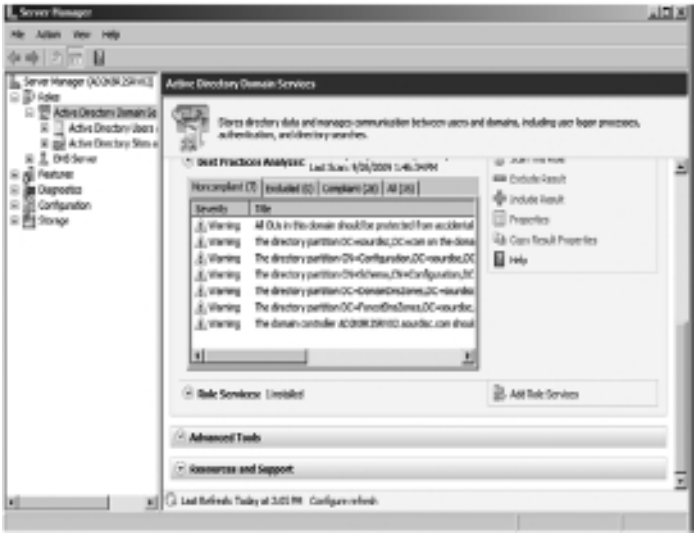

To open the Server Manager, select Start ➢ Programs ➢ All Programs ➢ Administrative Tools ➢ Server Manager. Using the Server Manager, you can expand the roles container and navigate to the Active Directory Domain Services role. Then scroll down to the Best Practices Analyzer section and scan the role using the ADBPA utility. In Figure 4.25, you can see that there various warnings that should be corrected as a result of the installation of Active Directory Domain Services on this virtual machine.

You will also use the Server Manager to perform the following tasks:

- Configure the PDC emulator to external stratum-1 time source
- Back up the Active Directory forest and domain partitions and system-state data
- Protect the Active Directory containers from accidental deletion
- Check and monitor replication to ensure that Active Directory is happy and healthy

Configure and Synchronize Time

We mentioned earlier in the chapter that there some differences in the way servers behave when they are physical vs. virtualized. Time synchronization is one of the biggest differences in this behavior. In most applications, the virtual machines consume CPU cycles and are therefore scheduled by the VMkernel. Active Directory presents a unique challenge in that most medium and smaller domain controllers don't require very much CPU resources to run optimally. What this means is there may be a lot of idle time on the virtualized domain controller. This is the challenge; idle virtual machines don't get scheduled by the VMkernel, and time synchronization will be virtually (pardon our pun) nonexistent. A virtual domain controller's time can drift significantly within a matter of hours. To prevent this from occurring, you must synchronize the PDC emulator with an external stratum-1 time source. Figure 4.26 shows an example error message in the Event Viewer when Active Directory is not synchronized or improperly synchronized. This will affect time in the whole forest.

FIGURE 4.26
This is an example error message when the Active Directory time is not synchronized.

The PDC emulator is responsible for the following in an Active Directory forest:

- It will act as the PDC for down-level clients such as Windows NT, Windows 95/98, and so on.
- The PDC emulator receives urgent replication of password changes within the domain. Because password changes take time to replicate to all domain controllers in an Active Directory domain, the PDC emulator will receive notification of password changes immediately and will accept logon requests from a client if they fail at another domain controller within that domain.
- The PDC emulator also serves as the machine that all domain controllers will use to synchronize their clocks. It is a Microsoft best practice to synchronize the PDC emulator with a reliable time source. We recommend that the PDC emulator be synchronized with an external stratum-1 Internet time source.

For the proof-of-concept example, we will show how to synchronize time using option 1 where we will create a Group Policy object (GPO) and query for the PDC emulator role in the forest. In the following text and figures, we will show you how to query for the PDC emulator and synchronize its time with an Internet stratum time source.

When creating the Group Policy to synchronize for the Active Directory domain controllers, you can navigate to the domain controller's organizational unit (OU) and create a new GPO, as shown in Figure 4.27. From there you edit the Group Policy object by right-clicking the object.

In the Group Policy management console, navigate to Computer Configuration ➢ Administrative Templates ➢ System ➢ Windows Time Services, and modify the parameters as shown in Table 4.2.

Once you have modified the required parameters listed in Table 4.2, you will then need to stop and restart the Windows time services on the domain controller. You accomplish this task by opening a command window and typing the following command at the prompt:

```
net stop w32time && net start w32time
```

This will stop and restart the Windows time services on the domain controller. Now that you have Group Policy configured, you can filter the policy so it applies only to the PDC emulator role in the Active Directory forest. You do this by navigating in the Group Policy management console to the Windows Management Instrumentation (WMI) Filters container and right-clicking to create a new WMI filter. Give the new filter a name and a description befitting the filter. After completing this step, click the Add button, and enter the following query:

```
Select * from Win32_ComputerSystem where DomainRole = 5
```

FIGURE 4.27
Creating a Group Policy object to synchronize time

Table 4.2: Windows Time Services Synchronization Parameters

Time Parameter	Policy Location
Announce Flags: Change from 10 to 5	Global Configuration Settings
NtpServer: timekeeper.isi.edu,0x1	Configure Windows NTP Client
Type: Change from NTDS5 to NTP	Configure Windows NTP Client
Enable the Windows NTP Client	Windows NTP Client
Enable the Windows NTP Server	Windows NTP Server

WMI uses different numbers to define the roles that servers play in the Active Directory forest. You can use these roles along with WMI filters to produce some action based on the computer roles. This can be a very effective method for obtaining predictable outcomes based on some action to be taken.

Table 4.3 lists the various server roles available in the forest.

Table 4.3: Computer System Roles

Number	Role
0	Stand-alone computer
1	Member workstation in domain
2	Stand-alone server
3	Member server in domain
4	Backup domain controller (Windows NT)
5	Primary domain controller

When you query for the domain role of 5, you are in fact querying for the PDC emulator role. By filtering for this role, you can ensure that the PDC emulator and only the PDC emulator will be synchronized with the stratum-1 Internet time source. We personally like this method of synchronization because it avoids editing the registry, and we know that some folks are not comfortable with editing the registry to modify the Windows time service parameters directly. Also, if you transfer the PDC emulator role to another domain controller, you will have to modify that domain controller's Windows time service parameters. Using a Group Policy object with a WMI filter to synchronize the PDC emulator eliminates the need to edit the domain controller's registry (Figure 4.28), and if you lose the domain controller the PDC emulator is running on or transfer/seize the PDC emulator role, the WMI filter will query for the role during the next Group Policy

refresh, find the new domain controller hosting this role, and apply the Group Policy of time synchronization to that domain controller.

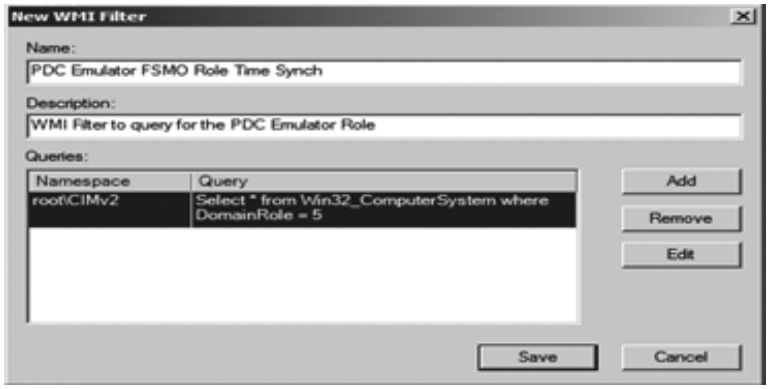

FIGURE 4.28
WMI filter to synchronize the PDC emulator to Internet stratum-1 time source

Once you have created the WMI filter, you must then link the timer server synchronization Group Policy that you created to the newly created PDC emulator FSMO Role Time Synch WMI filter. After you have completed this procedure, you will have successfully completed the time synchronization procedure. Figure 4.29 shows an event log message stating that the Windows Time service has successfully synchronized time with an external time source. From here on out, the PDC emulator will advertise itself as a reliable time source to the forest and the domain. Time services will then synchronized throughout the Active Directory forest.

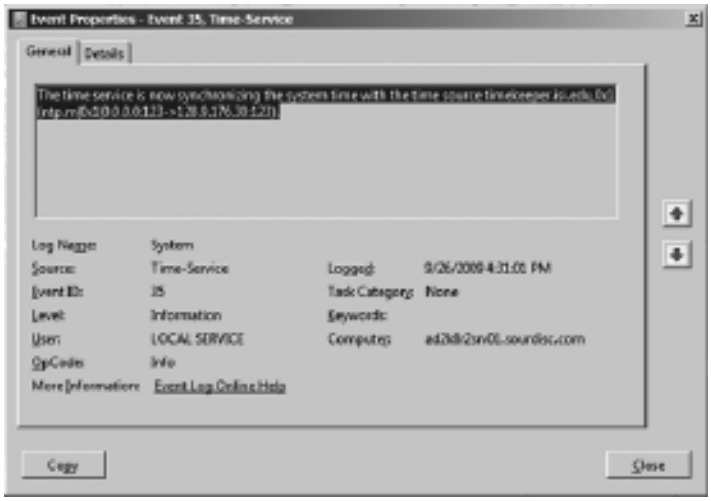

FIGURE 4.29
Viewing the system event log to determine whether the time is synchronized within the domain

Back Up Active Directory Domain Controllers

Now that you have installed and configured virtualized Active Directory domain controllers in the Source Discovery forest, you must follow the Microsoft and VMware best practices for virtualizing Active Directory by backing up the newly created domain controllers. We will not cover the restore procedure because there are numerous Microsoft technical articles about backup and

restore procedures for Active Directory. Our main focus here is to show you how to back up the domain controllers to ensure compliance with the Active Directory Best Practices Analyzer.

Install Backup Features

To install the Windows Server Backup features onto the domain controller, execute the following tasks:

1. Select Start ➢ Programs ➢ All Programs ➢ Administrative Tools ➢ Server Manager to open the Server Manager.
2. Select Features.
3. Click the Add Features link in the Features window.
4. In the Add Features Wizard, select the Windows Server Backup Features check box, and click Next.
5. Click Install to start the installation of the Windows Server Backup features onto the domain controller.

Configure the Domain Controller for Backups

Once the Windows system backup features are installed on the domain controller, you can proceed to configure the domain controller for automated backups, or you can perform a manual one-time backup. We will show how to perform a manual one-time backup to satisfy the ADBPA wizard requirements.

Expand the container if you want to enable the command-line features. Enabling the command-line features means you can use PowerShell commands and scripts in the Windows system backup. Figure 4.30 shows the installation of the Windows system backup features on the domain controller.

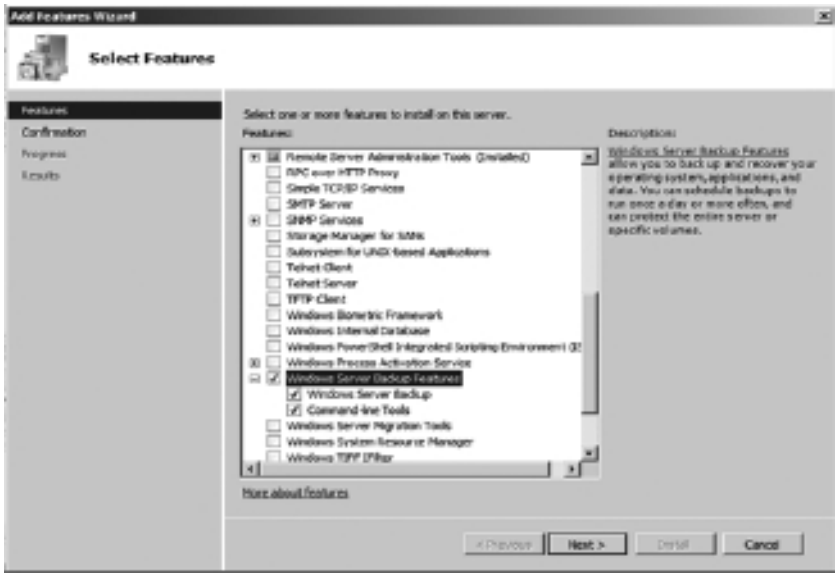

Figure 4.30 Using Add Features to install the Windows Server Backup utility

As a best practice, we recommend that you back up your domain controllers at a minimum at least once a week. The Windows Server Backup utility gives you multiple options for backing up your virtualized domain controllers. At a minimum, we recommend backing up the system state as well as all of the system drives. You can back these up separately or together. You can also create a system recovery disk that will enable you to do a bare-metal restore if necessary. We recommend also using this option. Figure 4.31 shows the backup options of scheduled or manual options for a customized backup.

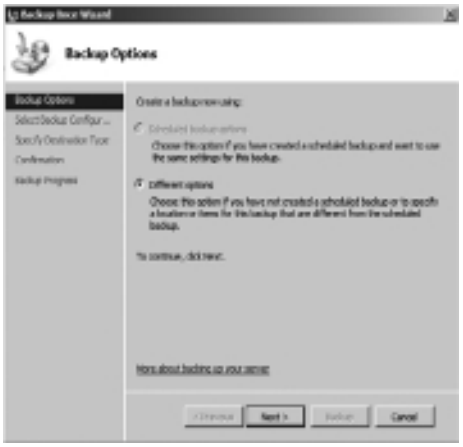

FIGURE 4.31
Selecting additional features for the backups

After selecting your backup options—which include system-state backups, individual drive backups, bare-metal backups, VSS full and copy backup options, and full server backups—you also must choose a destination for your backup. In our example, we used a network share to store our backup images. You have the option of using a network share, CD/DVD drive, or hard drive to store all of your backup images. Whatever medium you decide to use, just ensure there is ample space to store your backup images. The backup utility will utilize VSS to quiesce the Active Directory database, take snapshots, and back up those snapshots to the specified medium.

One thing that we should mention is that if you choose a network share for your backup destination, ensure that you enable write permissions for the group or user who you specify will have access to the network share. Once the backup has completed successfully, you should receive a successful completion notification on your domain controller (Figure 4.32). You should never keep a system-state or domain controller backup for more than the tombstone lifetime setting in Active Directory. For Windows Server 2003 SP1, in certain conditions the tombstone lifetime has been increased from 60 days to 180 days. For more information on tombstone lifetime expiration, see the Microsoft TechNet site.

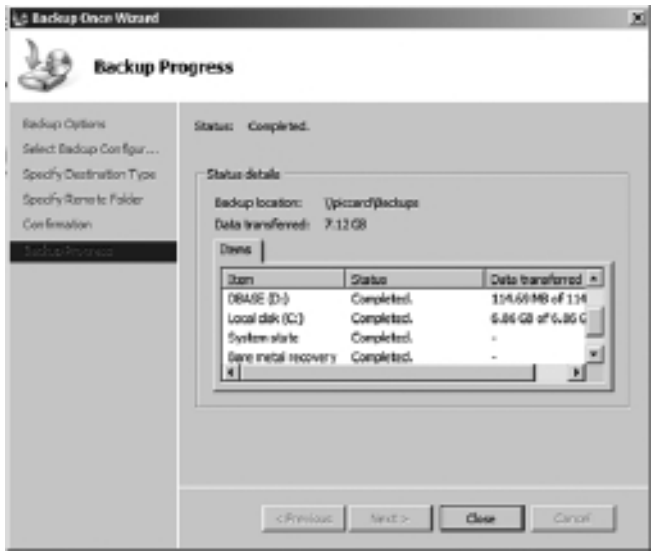

FIGURE 4.32
Backup process indicator showing successful completion of the manual backup

Prevent Active Directory Deletions

How many times have you been the victim of a system administrator making an honest mistake and accidently taking out 1 or 100 or 1,000 users at a time by deleting an organizational unit? Well, obviously this has been happening too much because Microsoft actually included a mechanism to prevent deletion, whether accidental or malicious, of the Active Directory organizational units and other containers. This is a welcome addition to Active Directory, but you must use it! The Active Directory Best Practices Analyzer ensures that you implement this when you scan the Active Directory Domain Services role (Figure 4.33). To be compliant with Microsoft best practices for Active Directory, you must implement this feature:

1. Open the Server Manager.
2. Expand the Active Directory Domain Services role, and then expand the Active Directory Users and Computers (ADUAC) container.
3. Select the sourdisc.com container, and select the view context at the top of the Server Manager.
4. Select Advanced Features. This will enable you access to the full properties on each object that you select. You will need to protect all of the objects while excluding the lost and found container.
5. Right-click the object, be it a built-in container or an organizational unit, and select the properties for the selected object.

6. Choose the Object tab. Select the Protect Object From Accidental Deletion box, and click the OK.

FIGURE 4.33
ADBPA message about protecting all Active Directory organizational units

And that's really all there is to this process. Again, as we stated earlier in this procedure, you must implement this feature, or every time that you scan the Active Directory Domain Services role using the ADBPA (Figure 4.34), you will receive warnings about your built-in containers and organizational units not being protected from accidental deletion. You will also leave yourself exposed to accidental and malicious deletion of critical components of the Active Directory infrastructure. Use this feature! It is probably one of the most useful new additions to Active Directory Domain Services for the Windows Server 2008 and R2 editions.

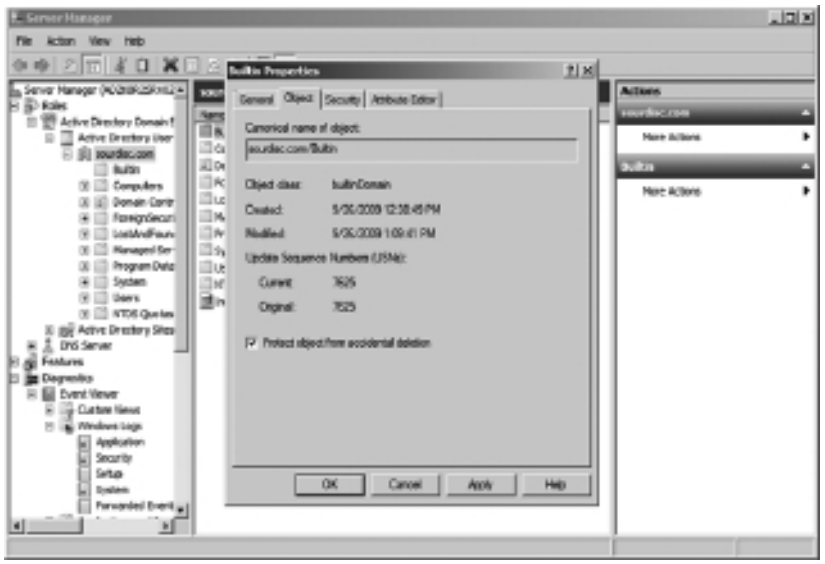

FIGURE 4.34
View of a built-in object protected from accidental deletion

Monitor Active Directory Replication

As we mentioned earlier in the chapter, proper replication is mandatory in ensuring that Active Directory is healthy and functioning properly. In this section, we'll cover how to use three tools to ensure the health of the Active Directory forest. As important as replication is to Active Directory health, this is usually one of the most neglected areas of operation.

RepAdmin

You can use the RepAdmin tool to view and verify that your domain controllers are replicating changes to their peers in the forest. The tool will tell you when the last successful replication happened, to which domain controller, over which transport type (RPC/SMTP) the replication happened, and the time and date that the replication took place. To run the replication administrator, open a command window, and type the following command:

```
Repadmin /showrepl
```

This command shows the various naming contexts and partitions of the domain controller and the last successful replication attempt on the domain controller the command is being executed on.

In Figure 4.35, the output shows how to use the `/showrepl parameter` command.

Select Administrator: Command Prompt
C:\>repadmin /showrepl
Repadmin: running command /showrepl against full DC localhost
Default-First-Site-Name\AD2K8R2SRV02
DSA Options: IS_GC
Site Options: (none)
DSA object GUID: b1a755e1-8552-459a-bb3d-6b81d27ef12a
DSA invocationID: 43132754-f983-4a5f-8da8-5acd30dcc693
==== INBOUND NEIGHBORS ======================================
DC=sourdise,DC=com
Default-First-Site-Name\AD2K8R2SRV03 via RPC
DSA object GUID: 125d2c61-59d4-4ada-8fa7-ff17ea87eab3
Last attempt @ 2009-09-28 10:51:01 was successful.
Default-First-Site-Name\AD2K8R2SRV01 via RPC
DSA object GUID: bb88d8eb-77e6-4257-8b21-bb883ef1b5e6
Last attempt @ 2009-09-28 10:51:01 was successful.
CN=Configuration,DC=sourdise,DC=com
Default-First-Site-Name\AD2K8R2SRV03 via RPC
DSA object GUID: 125d2c61-59d4-4ada-8fa7-ff17ea87eab3
Last attempt @ 2009-09-28 10:51:01 was successful.
Default-First-Site-Name\AD2K8R2SRV01 via RPC
DSA object GUID: bb88d8eb-77e6-4257-8b21-bb883ef1b5e6
Last attempt @ 2009-09-28 10:51:01 was successful.
CN=Schema,CN=Configuration,DC=sourdise,DC=com
Default-First-Site-Name\AD2K8R2SRV01 via RPC
DSA object GUID: bb88d8eb-77e6-4257-8b21-bb883ef1b5e6
Last attempt @ 2009-09-28 10:51:01 was successful.
Default-First-Site-Name\AD2K8R2SRV03 via RPC
DSA object GUID: 125d2c61-59d4-4ada-8fa7-ff17ea87eab3
Last attempt @ 2009-09-28 10:51:01 was successful.
DC=DomainDnsZones,DC=sourdise,DC=com
Default-First-Site-Name\AD2K8R2SRV01 via RPC
DSA object GUID: bb88d8eb-77e6-4257-8b21-bb883ef1b5e6
Last attempt @ 2009-09-28 10:51:01 was successful.
Default-First-Site-Name\AD2K8R2SRV03 via RPC
DSA object GUID: 125d2c61-59d4-4ada-8fa7-ff17ea87eab3
Last attempt @ 2009-09-28 10:51:01 was successful.
DC=ForestDnsZones,DC=sourdise,DC=com
Default-First-Site-Name\AD2K8R2SRV01 via RPC
DSA object GUID: bb88d8eb-77e6-4257-8b21-bb883ef1b5e6
Last attempt @ 2009-09-28 10:51:01 was successful.
Default-First-Site-Name\AD2K8R2SRV03 via RPC
DSA object GUID: 125d2c61-59d4-4ada-8fa7-ff17ea87eab3
Last attempt @ 2009-09-28 10:51:01 was successful.
C:\>

FIGURE 4.35
Output from the RepAdmin utility

We issued this command from domain controller AD2K8R2SRV02, and in the figure you can see that it has replicated changes successfully to its other partners AD2K8R2SRV01 and AD2K8SRV03 in the site. You can also see that it has successfully replicated the schema, configuration, and domain naming contexts as well as the forest and domain DNS partitions. The RepAdmin tool will

also allow you to produce simple reports on the overall status of your Active Directory replication. The following command will give you a detailed report including errors on the overall health status of your domain controllers within the Active Directory forest:

```
Repadmin /replsum /bysrc /bydest /sort:delta
```

As you can see in Figure 4.36, the largest delta or time elapsed since the last successful replication is only about 40 minutes. In the third column, you should notice that there have been no replication failures out of a total of five replication attempts and a zero-failure percentage along with no replication errors. As you can see, RepAdmin is a valuable tool for ensuring that your Active Directory forest is healthy and functioning efficiently.

Administrator: Command Prompt
C:\>md Temp
C:\>repadmin /replsum /bysrc /bydest /sort:delta > c:\Temp\RepAdmin_report.txt
C:\>repadmin /replsum /bysrc /bydest /sort:delta
repadmin /replsum /bysrc /bydest /sort:delta
Replication Summary Start Time: 2009-09-30 19:26:03
Beginning data collection for replication summary, this may take awhile:
.....
Source DSA largest delta fails/total %% error
Destination DSA largest delta fails/total %% error
C:\>

FIGURE 4.36
RepAdmin report on the virtual domain controllers in Source Discovery

REPLICATION MONITOR

Although ReplMon is not officially supported on Windows Server 2008 platforms, it does in fact work. This tool is the graphical user interface (GUI) version of RepAdmin. It can show you a graphical representation of the replication traffic in your Active Directory forest. It also allows you to view the naming contexts and partitions in more detail. When you leave this tool running and configure it to do automatic updates, it will monitor replication in real time and give you a viewable report in the GUI detailing the overall status of your domain controller replication.

You can see in Figure 4.37 that ReplMon is replicating all naming contexts and partitions properly and without any failures or errors. This tool also has an option to graphically view Active Directory domain controller replication traffic.

ACTIVE DIRECTORY BEST PRACTICES ANALYZER

When you have satisfied all of the ADBPA recommendations for virtualizing Active Directory domain controllers, you can run the ADBPA one final time. Figure 4.38 shows that we still have one warning message displayed by the ADBPA. After opening the Server Manager, we can expand the roles container and navigate to the Active Directory Domain Services role, scroll down to the Best Practices Analyzer section, and scan the role using the ADBPA utility.

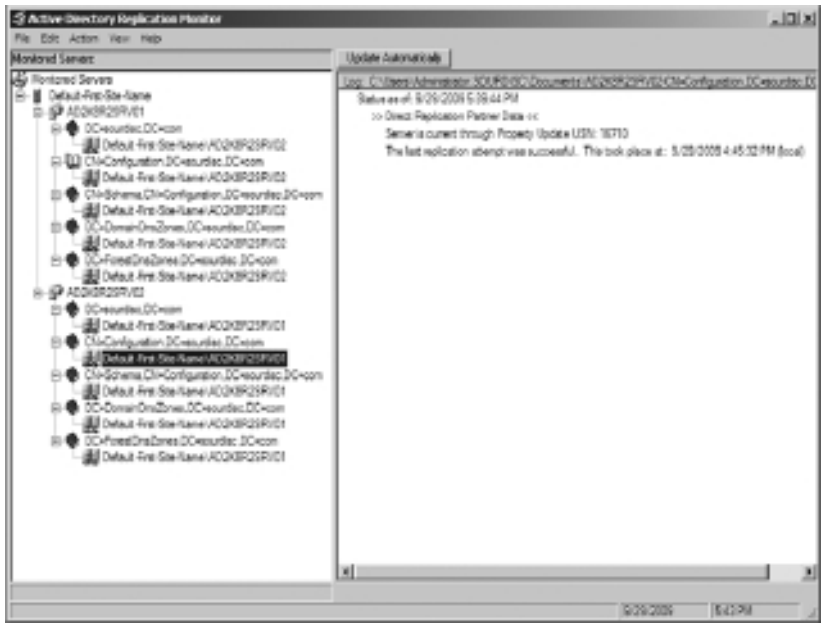

FIGURE 4.37
Using the ReplMon utility to monitor replication

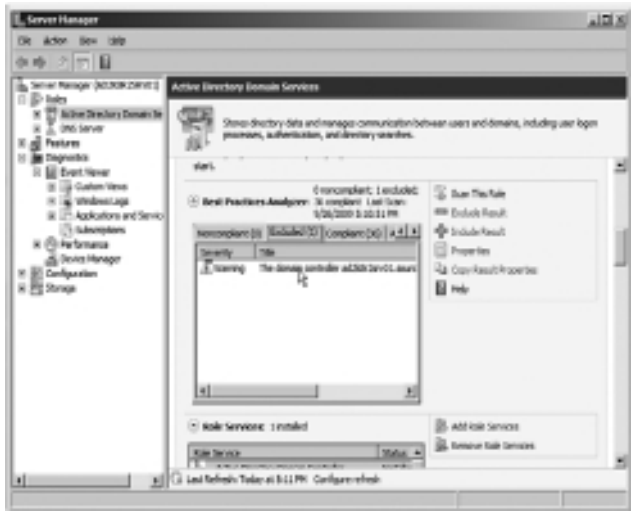

FIGURE 4.38
Active Directory Best Practices Analyzer scan results

In the results, we actually have 36 compliant results, 0 noncompliant results, and 1 excluded result. This is a warning message with recommendations for running Active Directory domain controllers virtualized (Figure 4.39). It says you should review all of the best practices for virtualizing Active Directory domain controllers. We have excluded this result, because it has no performance or negative effect on the virtualized domain controllers in the Source Discovery Active Directory forest. You can exclude this warning from the results for future scans.

FIGURE 4.39
Warning about virtualizing Active Directory domain controllers

NOTE For more information on this warning, please see the Microsoft TechNet article at `http://technet.microsoft.com/en-us/library/dd723681(WS.10).aspx`

Summary

In this chapter, we covered the considerations and best practices for virtualizing your Active Directory domain controllers. First we talked about the characterization of your existing Active Directory environment and how this plays an important part in your design considerations for your virtual design and implementation. Then we talked about the decision of whether to deploy 32-bit or 64-bit domain controllers in your environment. Caching larger directories in memory will require the 64-bit version of Windows Server and will yield significant performance benefits because you will not place extreme demands on the storage subsystem. We then talked about network configuration and how this can play an important in keeping your Active Directory healthy and happy.

We talked about the use of the latest virtual network interface card and teaming for your virtual switches that your Active Directory domain controllers are plugged into. We then touched on vSphere host placement and configuration for your Active Directory domain controllers. It is important to place domain controllers close to users where counts are above 25 users in a site to prevent authentication over your precious WAN links. Placing vSphere hosts near users can also assist in simplifying disaster recovery because virtual machines on failed hosts can be copied to functioning hosts and made operational. We then discussed the importance of time synchronization for your Active Directory domain controllers.

We talked briefly about the Kerberos protocol and its sensitivity to time synchronization. To ensure a happy and healthy Active Directory that allows users to authenticate for network resources, you must ensure your PDC emulator for the forest is synchronized to an external

stratum-1 time source. Once synchronized, your PDC emulator for other domains within your Active Directory forest will synchronize with the forest PDC emulator and provide accurate time for all members of the Active Directory forest/domain. We then moved on to cover Active Directory backup. Although it is important to ensure time synchronization and proper replication for a healthy and happy Active Directory, it is equally important to ensure that you have a well-defined and executed backup plan for Active Directory to assist you in properly restoring failed domain controllers back into service.

Improperly restoring an Active Directory domain controller can lead to serious domain and even forest corruption, ultimately requiring you to rebuild or redeploy Active Directory. We then talked about the importance of securing the virtual domain controller configuration files to ensure that unauthorized users do not gain access to information that they should not have and also to prevent accidental and malicious acts against virtual Active Directory domain controllers.

Lastly, we covered how to use two tools to monitor Active Directory replication. First we covered the RepAdmin tool and showed you how to monitor Active Directory replication as well as generate simplified reports of replication history for your domain controllers. Next we covered ReplMon, the GUI version of RepAdmin. This tool actually allows you to view the various Active Directory partitions and ensure that replication is occurring as expected. There is also an option for viewing the replication topology of your Active Directory forest. You can see which domain controllers are replicating with their replication partners.

Using these tools can ensure that your Active Directory forest/domains are healthy and processing logons, replicating changes and directory information in a timely and most efficient manner. Active Directory has been successfully virtualized by many of our customers and has been in service for extended periods without failure. Although your implementation will be different from our proof of concept, it is important to consider all of the best practices that we discussed in the chapter and implemented in the proof of concept to ensure your Active Directory virtualization implementation is a successful one.

Chapter 5

Virtualizing Internet Information Services

The release of Windows Server 2008 introduced a reworked version of IIS that included new features such as modular application design, tighter integration with ASP.NET, and configuration based on clear-text XML files. Virtualized web servers can take advantage of the new modular design to utilize today's multicore servers more efficiently. By installing only the required components, you can minimize the required virtual hardware and increase the capacity of your underlying physical hardware.

Windows Server 2008 R2 delivers an updated version of IIS to provide some of the features that didn't quite make it to the initial release. Officially identified as IIS 7.5, the updated version includes .NET support for Server Core and support for integrated WebDAV, among others.

IIS may not come across as a difficult-to-virtualize application, and it isn't. As with any application you choose to virtualize, if you know its characteristics, plan accordingly, and take advantage of resources available to you, you'll be quickly on your way to a smooth, virtualized IIS environment. That's exactly what we'll focus on in this chapter.

IIS can be a simple, unassuming application, one that we, as systems administrators, take for granted. This is probably the number-one reason why systems administrators classify IIS servers as "low-hanging fruit" when it comes to a virtualization project. Why is this? Well, chances are our experience with IIS has taken one of two forms:

A request comes in from our apps team for a web server running on the Windows platform. Simple—you click a few buttons, follow a few wizards, and IIS is installed. You hand the system over to the apps folks, and your job is done.

A new application has come across your desk. As any good sys admin does, you review the requirements (before you do anything else, right?), see that the application requires IIS, and as in the previous scenario click a few buttons and are good to go. You install the app, which takes care of creating the website, making its necessary changes, and you're done.

This leaves you with a false sense of knowing IIS. Sure, you've deployed IIS and may have poked around, but what do you really know about it? Do you have any idea what type of load serving web pages puts on the system? Do you have any idea if you're serving static or dynamic content? Are you CPU or memory bound? What about disk I/O? Well, now that we've gotten the gears turning, we'll kick the IIS tires and see what it is you must consider when deploying IIS in a virtualized environment.

Deploying Internet Information Services

Before you begin deploying IIS servers, you need to plan the environment. Luckily, working with a vSphere infrastructure usually means your hardware has been established and you will simply be carving up resources for your virtual machines. How you build these virtual machines will depend on factors such as company standards, licensing, and technical requirements.

Because IIS has been part of the Windows operating system since Windows Server 2000, you don't have to install a specific version of Windows to support IIS; rather, you can make your OS choice based on other factors. Which OS is right for you? Chapter 3 did a good job of breaking down the various versions of Windows Server 2008 supported on the vSphere platform, but how they differ might steer you toward one or the other.

Choosing a Flavor of Windows Server 2008

As you know, Windows 2008 comes in a few different flavors: Standard, Web, Enterprise, and Datacenter. Most of us are familiar with the Standard and Enterprise editions of Windows Server. The major difference here is the memory support: 4 GB vs. 64 GB for x86 versions and 32 GB vs. 2 TB for 64-bit versions. There is also support for failover clusters in the Enterprise version.

The Web and Datacenter editions of Windows server are usually left for specialized installations. Web edition is pretty self explanatory; this version of Windows Server is able to run IIS 7.0, ASP.NET, and SQL Server to support local applications. Windows Server's Web edition will not run other applications such as Microsoft Exchange Server. Datacenter edition is usually deployed in large-scale SQL installations because of the large processor (up to 32 processors in the x86 version and 64 processors in the 64-bit version) and memory support (64 GB in the x86 version and 2 TB in the 64-bit version).

Finally, you have the option to install the Server Core versions of Standard, Enterprise, and Datacenter as a host for your IIS installation. Server Core is a bare-bones installation of Windows Server 2008 and is especially useful for web servers because the minimal installation reduces the attack surface. Server Core installs with no Explorer shell, which means management is handled at the command line when on the console or remotely using Remote Desktop Services or the IIS Manager.

NOTE If you plan to use the Windows Server 2008 RTM version (which means you are installing IIS 7.0) on Server Core, ASP.NET cannot be installed. This restriction has been lifted in IIS 7.5, which is installed with Windows Server 2008 R2.

The hardware limitations of the various versions of Windows Server 2008 make choosing the correct version fairly trivial once you have chosen the hardware makeup of each virtual machine. If the requirement is that the VM be built with eight CPU cores and 16 GB of memory, then deploying Windows 2008 Standard edition is not an option, and you need to look at the Enterprise or Datacenter edition. Deciding the amount of hardware, or virtual hardware, required is a critical step in the planning phase. The price difference between the various flavors of Windows Server licenses could be enough to sway decisions one way or another.

NOTE Later in the chapter, we will explore the various ways to monitor compute resource utilization on an IIS server, which will help you determine your bottlenecks and where to add resources.

Once the virtual hardware has been sized and the OS is installed, adding the role of IIS follows the same process for the various editions of Windows Server 2008. The process for installing the IIS role is described in the following sections.

Installing IIS

With the new modular design of IIS 7, you are able to gain security and performance by only installing the modules you require to serve the content necessary. The default installation of IIS 7 includes the role services necessary to serve and compress static content. Without installing any additional modules, you can make minor customizations such as editing default documents and HTTP errors, as well as logging and monitoring server activity. Before installing any roles, you want to be sure that the Administrator account has a strong password, that you have configured and confirmed all network settings, and finally that you have updated your OS with all the latest security patches using Windows Update

IIS Manager

The IIS Manager application has been completely revamped to be more user-friendly and intuitive. The default installation of IIS provides basic management through the GUI. As you install additional IIS features, the IIS Manager is extended with the tools required to manage these features. IIS Manager also provides the ability for third-party vendors or systems administrators to extend IIS Manager with their own tools.

IIS Manager is the recommended management interface for all IIS installations. This especially includes IIS installations on the Windows 2008 Server Core operating system. The lack of a GUI interface in Server Core means that using IIS Manager requires the connection to be made remotely. The ability to remotely manage IIS is not installed with the default IIS installation. The following procedures describe how to enable remote management on the Server Core version of Windows Server 2008 installations. In the proof of concept later in this chapter, we will walk you through enabling remote management in the full version of Windows Server 2008 since Source Discovery does not currently use Server Core installations.

Remote Management in Server Core

To enable remote management in Server Core, perform the following steps:

1. Install the IIS Remote Management service using the following command:

   ```
   start /w ocsetup IIS-ManagementService
   ```

2. Enable the IIS Remote Management service by issuing the following command:

   ```
   reg add HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\WebManagement\Server /v
   EnableRemoteManagement /t REG_DWORD /d 1
   ```

3. Start the management service using the following command:

   ```
   Net start wmsvc
   ```

4. Set the Web Management Service to start automatically when the server boots up using the following command (the space between the = and `auto`, is required):

   ```
   sc config WMSVC start= auto
   ```

Command-Line Administration

For those of us who prefer a command-line interface or would like to script a set of repetitive tasks, the `appcmd.exe` tool is a welcome feature. Available for use in both full and Server Core installations, the `appcmd.exe` tool allows for adding, deleting, and modifying websites and application pools, starting and stopping websites and application pools, and viewing and modifying the configuration of both IIS and ASP.NET. The following is the syntax, as shown in the command help, accessible by entering the command:

```
Appcmd.exe /?
```

You can get detailed help for each object by entering the following command (shown in Figure 5.1):

```
Appcmd.exe [object] /?
```

C:\Windows\System32\inetsrv>appcmd.exe
General purpose IIS command line administration tool.
APPCMD (command) (object type) (identifier) (/parameter1:value1 ...)
Supported object types:
SITE Administration of virtual sites
APP Administration of applications
VDIR Administration of virtual directories
APPPOOL Administration of application pools
CONFIG Administration of general configuration sections
WP Administration of worker processes
REQUEST Administration of HTTP requests
MODULE Administration of server modules
BACKUP Administration of server configuration backups
TRACE Working with failed request trace logs
(To list commands supported by each object use /?. e.g. 'appcmd.exe site /?')

FIGURE 5.1
Appcmd.exe
help menu

Fortunately, you can install IIS and create a template for future use. You learned how to build and deploy virtual machine templates in Chapter 3. With templates, you have a foundation that you can use to build your web farm environment. How you outfit your virtual machines will depend on the resources your web apps require. In the next section, you will look at the performance metrics that are important when planning the IIS server environment.

IIS Performance

Unlike other applications where you may have specific performance characteristics, the resources used by IIS will vary greatly depending on how the website or application is developed. In the absence of clear hardware requirements, we recommend one of the following two options:

- Use the Windows 2008 Server minimum hardware requirements at `www.microsoft.com/windowsserver2008/en/us/WS08-system-requirements.aspx` as the base, and tune the resources based on your workload.
- Use generally recommended settings based on the IIS 7.0 Administrator's Pocket Consultant publication from Microsoft Press at `http://technet.microsoft.com/en-us/library/cc268240.aspx` as the base, and tune the resources based on your workload.

Databases tend to perform best when given memory for caching and more disk spindles to speed up reads and writes. IIS applications can vary from being CPU, memory, or disk intensive

in nature. Unfortunately, there isn't an easy way to cookie-cut a high-performing IIS installation. Because of this, you will need to do some proactive work once your IIS application is installed and configured.

To properly pinpoint where your bottlenecks may lie or characterize the workload, you can utilize tools such as performance monitor. Performance monitor can be an overwhelming tool to use. Fortunately, you can concentrate on a few counters that will give you good insight into your system performance. Once you've characterized the performance of your application, you will need to make a few decisions on how to get the best performance out of your system.

Depending on the application being hosted on a particular system, you usually have the option to either scale up or scale out.

Scaling up This refers to adding resources to the individual system so that you can process more requests from a single system. In a virtual environment, much like in the physical world, there is a point of diminishing return, where continuing to scale up becomes less cost efficient than scaling out. In the physical world, this usually occurs when all CPU sockets have been consumed by the most cost-efficient processor and the memory slots have been filled with the most cost-efficient memory modules. In the virtual world, you can think of the server's resources as a pool of resources to be shared by many virtual machines. Chances are you won't deploy a single virtual machine on a piece of hardware, which would hardly be efficient use of the physical server resources. This is when you would start considering a scaled-out architecture.

Scaling out When scaling out, you will deploy multiple systems that are configured identically to either handle more requests or host more websites. Deploying a scaled-out environment does come with more overhead. There is more time and effort required when configuring and maintaining a scaled-out environment; however, the benefits come in the ability to handle more requests with less resource contention and the ability to balance the workload among multiple systems.

SpecWeb2005 Results

Although not based on IIS when talking about scaling out, there's no better proof than the record-breaking results set by a single ESX host using the SpecWeb2005 workload. These results were the highest recorded on a 16-core system. The test used the scaled-out approach using multiple one-vCPU/8 GB memory virtual machines on a single host using the Banking, Ecommerce, and Support workloads that represent the workloads common for web servers. For more information and the full disclosure report, visit the SPEC website: `www.spec.org/osg/web2005`.

Whether you choose to deploy a single virtual machine with the maximum virtual hardware configuration or many smaller VMs in a scaled-out fashion, you will need to monitor the system's performance. Monitoring is done not only at the VM guest level using tools such as Perfmon but also at the ESX host level using native tools such as `esxtop` or vCenter performance graphs to identify any bottlenecks at the physical layer. In the following sections, you will learn which performance counters are relevant for IIS and ESX and the tools used to monitor those counters.

Monitoring Performance

Anyone who has launched the Perfmon tool in Windows knows that it can be quite overwhelming. Knowing what counters to monitor takes a bit of knowledge, and even when you know what to log, do you really know what is good and what is bad? In the following sections, we'll simplify things by taking a brief look at some of the most important counters to monitor on an IIS web server.

Processor Utilization

Processor bottlenecks in IIS will result in poor web page performance, slow loads, or even blank pages being displayed. High processor utilization can manifest from heavy ASP.NET code or simply too high a load for the IIS server. You can monitor the counters listed in Table 5.1 to determine whether you are experiencing a CPU bottleneck.

Table 5.1: Processor Counters

Object	Counters	Description
Processor	% Processor Time	The percentage of time the processor is busy. The value of this counter is determined by subtracting the percentage of time the idle process is running from 100 percent. Generally sustained values greater than 70 percent will indicate a processor bottleneck.
System	Processor Queue Length	The number of threads ready to execute on a processor but unable to do so because of a different active thread currently consuming the processing resources. A sustained queue of two or more threads while CPU usage is below 90 percent indicates a CPU bottleneck.
Active Server Pages	Request Wait Time	Time in milliseconds the most recent request was waiting in the queue.
	Requests Queued	Number of requests waiting to be processed.
	Requests Rejected	Requests that were discarded due to resources being unavailable.
	Requests/sec	Number of requests executed per second.
		Generally, request should not be queued, but if they are, the wait time should be low. Additionally, if requests are being rejected due to insufficient resources, you may be experiencing a CPU bottleneck.

Memory Utilization

Lack of memory can often be the leading issue of performance problems. High memory utilization can lead to excessive paging, which will result in high disk I/O and increased CPU utilization. Use the counters in Table 5.2 to help determine whether you are experiencing a memory bottleneck.

TABLE 5.2: Memory, Caching, and Virtual Memory Counters

OBJECT	COUNTERS	DESCRIPTION
Memory	Available Kbytes	Amount of physical memory available for use by processes on the server. Available memory should be no less than 5 percent of the total physical memory.
	Committed Bytes	Amount of committed virtual memory. Committed memory should generally be no more than 75 percent of the total physical memory.
	Cache Bytes	Total size of file system cache.
	Pool Paged Bytes	Paged pool is the area of system memory for objects that can be written to disk when not in use. Large page pool size relative to the amount of physical memory may indicate the need for more physical memory.
	Pool Nonpaged Bytes	The nonpaged pool is the area of system memory for objects that cannot be written to disk. Large nonpaged pool size relative to the amount of virtual memory allocated to the system may indicate an increase to the amount of virtual memory is required.
Internet Information Services Global	Current File Cache Memory Usage	Memory in use by the IIS file cache.
	File Cache Hits %	Ratio of cache hits to total cache requests. Sites with mostly static content should have high cache hit ratio, above 75 percent.
	File Cache Flushes	Indicates how quickly IIS is flushing files out of cache. Quick flush rates may indicate the need for longer time-to-live values for cached objects (ObjectCacheTTL). Flushes occurring too slowly may waste memory and indicate the need to decrease the time-to-live for cached objects.

DISK INPUT/OUTPUT

Chances are if you have properly sized the memory available to your IIS server, you will rarely see issues in the disk subsystem. You may notice degradation in performance if there is sufficient paging or if you are coexisting IIS with other more disk I/O-intensive applications. Use the counters in Table 5.3 to gauge the performance of your disk subsystem.

TABLE 5.3: Disk I/O Performance Counters

OBJECT	COUNTERS	DESCRIPTION
PhysicalDisk	% Disk Time	Monitor % Disk Time in conjunction with Processor\% Processor Time and Network Interface Connection\Bytes total/sec. High % disk time and relatively low processor and network utilization may indicate a disk bottleneck.
	Disk Writes/sec Disk Reads/sec Avg. Disk Write Queue Length Avg. Disk Read Queue Length Current Disk Queue Length	Disk read and writes per second indicate how much I/O activity is being generated on the disks. Read and write queue lengths tell you how many requests are waiting to be processed. In general, there should be few requests pending processing. Request delays are proportional to the length of the queues minus the number of drives in a RAID set.

NETWORK BANDWIDTH

Generally speaking, the bandwidth available to your web server will be far greater than the traffic that is being serviced by your website. Most network interface cards these days are capable of handling network speeds of up to 1 Gbps. This bandwidth is usually shared between the applications, management tasks such as remote desktop connections, and backups.

To monitor network bandwidth usage, employ the Bytes Total/sec counter. It monitors the bytes sent and received per second over each network adapter. Under typical load, this value should be 50 percent or less; otherwise, the system may experience issues during peak times.

All of the previously listed counters will provide valuable insight into how the guest believes it is performing. It is important to understand that the guest is abstracted from the fact that it is not running on dedicated hardware. This means that you may not be seeing the whole picture, and where the bottleneck lies may be hidden from these results. The following additional tools will help in gathering additional performance data.

VM STATISTICS IN PERFMON

After installing VMware tools in your virtual machine guest, two performance counters are installed, VM Memory and VM Processor. These counter objects use the VMStatsProvider to expose the performance for the VM, as seen by the ESX host. It is a little known fact that the higher utilized a VM or ESX host gets, the less accurate the in-guest performance-monitoring tools become. These in-guest tools provide the most accurate in-guest readings for processor and memory.

To view the available counters, follow these steps:

1. Launch `Perfmon.exe`.
2. Right-click the graph area, and select Add Counters.

3. Select either the VM Memory or VM Processor performance object, and select the desired counters (Figure 5.2 and Figure 5.3).

FIGURE 5.2
VM Memory performance counters

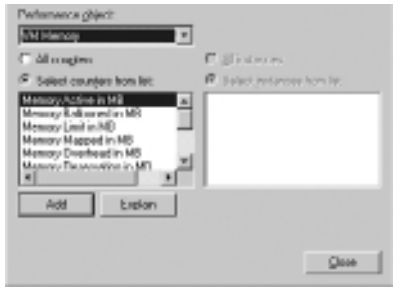

FIGURE 5.3
VM Processor performance counters

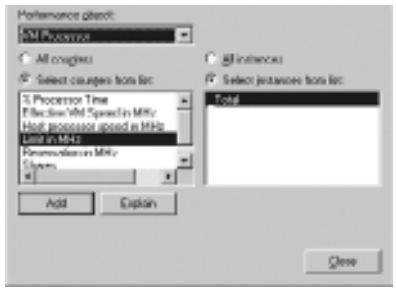

ESXTOP

The `esxtop` command-line tool provides a real-time view of the ESX server workload. To launch esxtop, you must SSH or open a console to the ESX host (Figure 5.4). In interactive mode, each resource is displayed on its own screen. Using the keyboard, you can change the resource that is being monitored and how much data for that resource is displayed.

FIGURE 5.4
esxtop

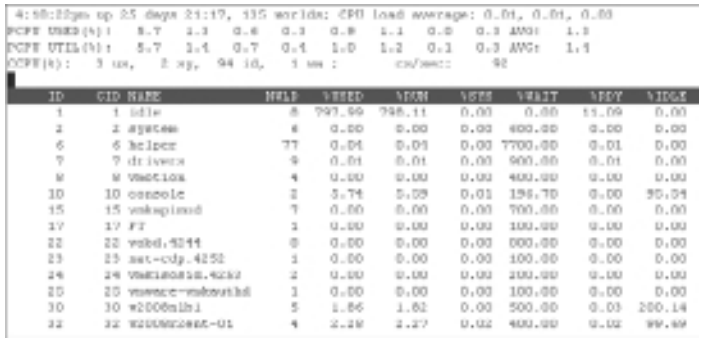

To toggle between the resource you are monitoring after launching `esxtop`, press the key that corresponds to the desired resource, as shown in Table 5.4.

TABLE 5.4: `esxtop` Resource-Monitoring Keys

RESOURCE	KEY
Processor	C
Memory	M
Network	N
Disk	D
ESX Disk	U
Virtual Machine Disk	V

The default data that is displayed for each resource will help provide details as to how the ESX host and the virtual machines are running and how many system resources they are using. If necessary, you can use the F key to modify the default columns shown in the display.

Using Perfmon and `esxtop`, you can track various aspects of performance from the IIS installation, the guest, and the underlying hypervisor. With the performance information in hand, you will be able to more precisely tune your virtual hardware to satisfy the requirements.

VMware AppSpeed

Any critical IIS installation will most likely be the point of entry to a larger enterprise system. Today most large enterprise systems are comprised of a web layer (that is, IIS), an application layer (that is, SharePoint), and a database layer (that is, SQL). With these layered application stacks, it can be extremely difficult to determine where a problem may lie when your customers call to complain about slow performance. The tools you were introduced to in the previous section are great for looking at the local system and even the underlying hypervisor, but what if that checks out?

VMware AppSpeed provides a deep view into the transactions occurring between virtual machines in your vSphere environment. By dynamically discovering applications currently running in your environment, AppSpeed can help map application and infrastructure dependencies. This data will translate into improved root cause analysis and decreased time to resolution in the future by helping you determine where in the stack the performance degradation may be occurring. In a physical environment, this type of data can be obtained only by using network-sniffing software or devices, span ports, and the assistance of the networking and possibly other teams. AppSpeed brings all the information available using the traditional tools plus application discovery and integration with vSphere.

Typically administrators monitor resource utilization and use that as a proxy for application performance. However, most application SLAs are defined from the end user's perspective. AppSpeed bridges this gap by providing application performance in terms of latency and transaction throughput. AppSpeed can then pinpoint the specific application transactions that are slow or most hit and their dependencies on other components in the application stack.

Installing VMware AppSpeed

VMware AppSpeed is an additional component of the vCenter family of management products. When you download AppSpeed, you will receive a ZIP file containing three files, as shown in Figure 5.5. The files will include an Open Virtualization Format (OVF) file that contains the packaging and distribution information required to properly deploy a virtual machine appliance, along with two virtual machine disk files (VMDK). Together these files will be imported into your vCenter server to create an AppSpeed collector and probes. Unzip these files to an easy-to-remember location.

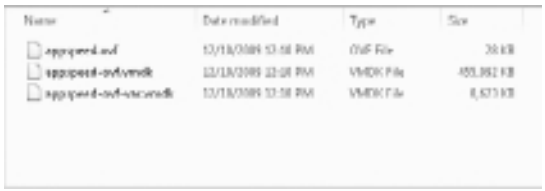

FIGURE 5.5
Contents of AppSpeed ZIP file

AppSpeed can provide valuable information about the health of your virtualized IIS environment. This is data that in a physical deployment would require cooperation between the network administrators and possibly hardware devices to aggregate and correlate the information collected. Virtualizing your IIS environment along with the other components in your multitiered application stacks allows for the use of tools such as AppSpeed to seamlessly manage your vSphere infrastructure. VMware AppSpeed is deployed during the proof-of-concept portion of this chapter to monitor Source Discovery's virtualized IIS environment.

To learn more about how you can take advantage of the features AppSpeed can provide, visit the AppSpeed page at `www.vmware.com/products/vcenter-appspeed/`.

Performance Tuning

Maximizing the performance of your IIS server and the application and websites themselves will ensure visitors to your websites will have the experience you expect. Using what you've learned from the previous section, you can gain knowledge into where performance bottlenecks may exist. Code that your websites depend on must also be finely tuned and developed as efficiently as possible to ensure you are performing as optimally as possible. Although this is usually an area we, as systems administrators, leave to the developers, you can take your knowledge of optimized resource utilization to point out anomalies that may help developers track down performance bugs in their code.

Besides providing the underlying computing resources for the IIS implementation system, administrators can assist in helping achieve the best performance possible, not only at the server level but at the IIS and website levels as well. In Chapter 3 you looked at how you could take a cookie-cutter approach to deploying high-performing Windows Server 2008 virtual machines. With the groundwork for server-level performance laid, you can concentrate on how to get the most out of IIS and the websites you intend to host.

IIS Application Tuning

When tuning the IIS application, the biggest bang for the buck comes in optimizing the `http.sys` HTTP listener. `http.sys` is the kernel mode device driver of IIS that helps improve performance by running with this elevated level of access to hardware and system resources. Introduced in IIS 6.0, `http.sys` helps improve performance by utilizing kernel mode caching. This allows requests for cached responses to be served without switching to user mode. Kernel mode request queuing is another benefit of running the listener outside of user mode. The elevated access given to `http.sys` allows requests to be forwarded directly to the appropriate worker process; if no worker process is available, the request is queued until a worker process becomes available.

By looking at the benefits of running the HTTP listener in kernel mode, it becomes pretty obvious that if you can tune `http.sys` to maximize its resource use, you can optimize IIS at the server level for all websites. Tuning how `http.sys` receives and responds to requests is handled in the HTTP service registry key:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\HTTP\Parameters
```

Table 5.5 lists the values used to control `http.sys` I/O processing.

Table 5.5: `http.sys` Registry Values

Registry Value	Description
UriEnableCache	This is a nonzero value that enables kernel mode response cache. When enabled, receive requests are served from kernel mode cache without transitioning to user mode. The default is on (1).
UriMaxCacheMegabyteCount	This is the amount of memory to be used for kernel mode response cache. If set to zero (0), the size is managed by the operating system.
UriMaxUriBytes	This determines the maximum size allowed for a response to be cached in the Uri cache. Items exceeding this value will not be cached. This value may be increased to take advantage of systems with larger amounts of physical memory. The default value is 256 KB.
UriScavengerPeriod	This is the time in seconds when the scavenger process will remove responses or fragments that have not been accessed in at least the UriScavengerPeriod. The default is 120 seconds.

Optimizing Websites

Now that we've gone through and covered how you can optimize the server as a whole, you can move on to working with individual websites. The characteristics of some websites might require more fine-grained tuning. The following optimizations can be performed per website as opposed to at just a server level.

HTTP Keep-Alives

HTTP keep-alive responses improve performance when there are multiple elements on a web page. By keeping client/server connections open across multiple server requests, new connections are kept to a minimum. The HTTP keep-alive response header is enabled by default.

Content Expiration

Content expiration marks each web page with an expiration time that the client browser will use to determine whether the page needs to be refreshed from the web server. This improves performance by reducing the frequency of client browsers requesting page refreshes.

To review HTTP keep-alive and content expiration settings, perform the following steps:

1. Launch IIS Manager.
2. Select the server object on the left pane, and double-click the HTTP Response Headers feature (Figure 5.6).

FIGURE 5.6
Open the HTTP Response Headers feature to edit settings.

3. In the Actions pane, select the Set Common Headers link (Figure 5.7).

FIGURE 5.7
Edit the HTTP Header settings by clicking the Set Common Headers link.

4. Review the settings, or make adjustments to the default settings (Figure 5.8).

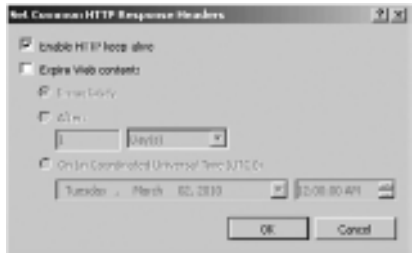

FIGURE 5.8
Keep-alives, enabled by default, enable web content expiration if desired.

Bandwidth Throttling

Bandwidth throttling can be used to limit the amount of network resource available to all websites or a specific set of websites. The default setting for bandwidth throttling is disabled. To review the bandwidth throttling settings, perform the following steps:

1. From within IIS Manager, expand the Sites node under the web server node in the left pane, and select the site to administer.
2. With the site selected, the Actions pane, on the right side of the IIS Manager window, provides links to the various actions available to each site. Select the Limits link under Configure (Figure 5.9), within the Edit Web Site Limits dialog box edit the bandwidth usage and connection timeouts (Figure 5.10).

FIGURE 5.9
Access the Edit Web Site Limits dialog box using the Limits link on the Action pane.

FIGURE 5.10
Edit the bandwidth usage and connection timeouts.

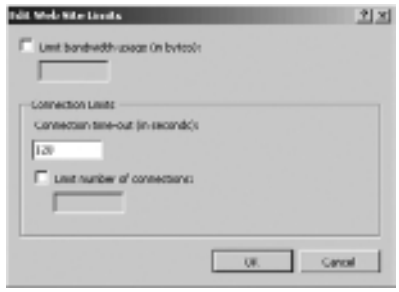

Output Caching

Output caching improves performance by storing a copy of a requested web page in memory. During subsequent requests for the same content, the response is served from memory as opposed to being regenerated. To access the Edit Output Cache Settings dialog box, perform the following steps:

1. From within IIS Manager, expand the Sites node under the web server node in the left pane, and select the site to administer.
2. With the site selected, double-click the Output Caching feature from the center pane.
3. Click the Edit Feature Settings link on the Actions pane (Figure 5.11) and enable or disable the cache settings (Figure 5.12).

FIGURE 5.11
Access the Edit Output Cache Settings dialog box using the Edit Features Settings link.

FIGURE 5.12
Enable or disable cache settings and set the cache size.

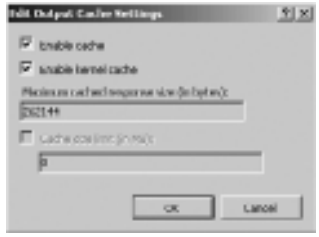

HTTP Compression

HTTP compression can reduce the amount of bandwidth used by static and dynamic content on your website and allow for more efficient use of your network resources. Dynamic content compression does require more CPU resources than static compression and should be limited

to servers that have slow network connections and CPU resources to spare. To edit the compression settings, follow these steps:

1. From within IIS Manager, either select the web server node to make server-wide changes or expand the Sites node under the web server in the left pane and select a specific site to administer.
2. With the server or site selected, double-click the Compression feature from the center pane.

NOTE By default, only the static content compression role service is installed with the default role services selected. To enable dynamic content compression, you must install the role service Dynamic Content Compression.

The center pane shows the options available. Options are different based on whether the focus is at the server level or website level. Figure 5.13 and Figure 5.14 show the different options.

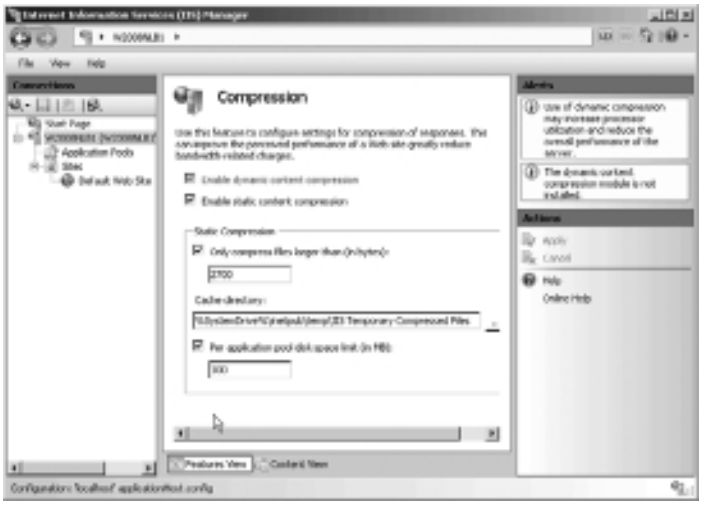

FIGURE 5.13 Server-specific compression settings

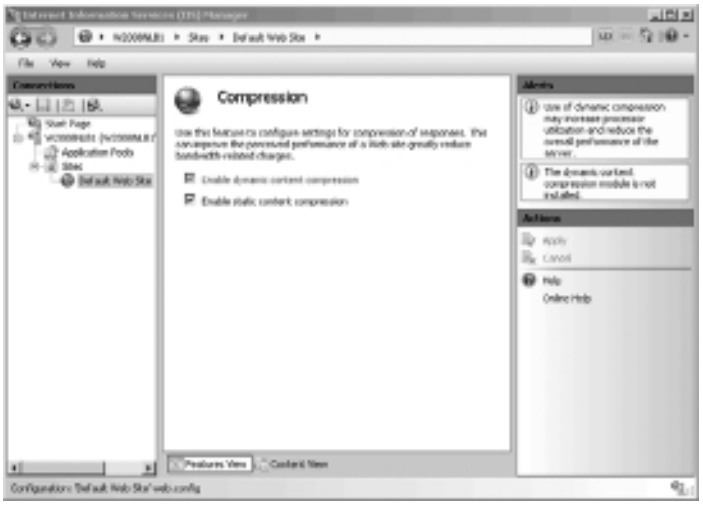

FIGURE 5.14 Site-specific compression settings

Becoming familiar with the previous settings will allow you to test different settings and optimize the responsiveness of your websites and servers. Once you have properly characterized your workloads and sized your virtual hardware accordingly, you can begin looking at what other optimizations may be possible. If the scaling-up approach becomes cost restrictive or you simply want to add some level of redundancy to your environment, clustering or load balancing may be the next feature to look into. In the next section, we cover a few methods of load balancing and what to look out for when load balancing in a vSphere environment.

Network Load Balancing

The web server, in this case the IIS server, is an extremely flexible application to deploy en mass. Where many of today's application servers must be scaled upward, web server front ends allow for outward or horizontal scaling. What this really means is where you would normally have to dedicate many of the resources available on a single ESX host to a single virtual machine, you can now spread those resources to smaller, leaner, more portable VMs. Why would you want to have more virtual machines to manage? Well, having more virtual machines means more partitioning of the underlying hardware, which means not only will you be able to utilize the ESX hosts' resources more efficiently, but you will also be adding significant levels of high availability.

DNS Round-Robin

In its simplest form, DNS round-robin provides load distribution rather than load balancing. This distribution is performed by randomizing the results of a DNS query for a specific host record. This randomization results in a different IP address being returned at the top of each query. Clients will generally use the first IP address returned by the query.

As an example, you can take the fictitious company Source Discovery whose website is at `www.sourdisc.com`. When you perform an `nslookup` on `www.sourdisc.com`, you see that the host record for `www` returns multiple IP addresses. On subsequent queries, the results will be returned in a different order than the previous queries (Figure 5.15).

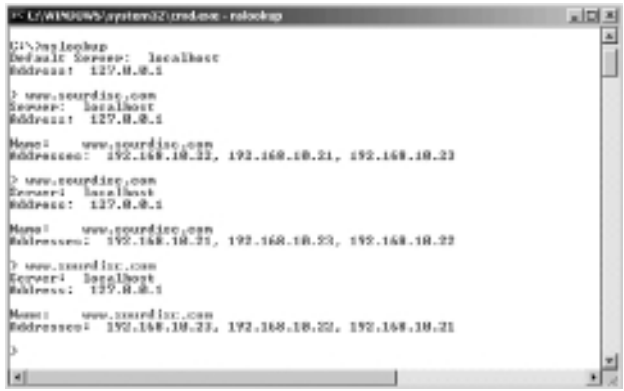

Figure 5.15
DNS round-robin returns each name request with a randomized list of IP addresses.

RFC3484: Default Address Selection for Internet Protocol Version 6 (IPv6)

Newer operating systems may react differently than expected when presented with a randomized list from DNS. These will mainly be OSs that were designed with IPv6 in mind. RFC3484 was written for compliance by all IPv6 nodes and introduces a few changes to the way a destination address is selected by the client. When the list of randomized addresses is returned to the client, the client will then reorder them based on a set of rules outlined in the RFC. A major contributor here is the rule that requires using the longest-matching prefix. What this means is that these clients will not simply take the first IP address in the query. Instead, they will work through an algorithm to make sure they attempt to connect to the closest IP address first. For more information, refer to the RFC page: `www.ietf.org/rfc/rfc3484.txt`.

Software Load Balancing

Since Windows 2000, the server family operating systems have shipped with the ability to provide a software load-balancing solution. Microsoft's Network Load Balancing (NLB) provides an out-of-the-box solution for simple load balancing. NLB runs as part of the operating system on all hosts in a cluster or pool and provides a distributed algorithm to balance network traffic among the cluster members. Features such as server failure detection and client affinity make NLB a favored load-balancing option over round-robin DNS. NLB is also the core to load balancing when using Microsoft Internet Security and Acceleration (ISA) server.

Virtual appliances have proven to be a good way for evaluating a feature set provided by a vendor of hardware appliances. On the VMware Virtual Appliance Marketplace, there are a number of virtual load-balancing appliances from various vendors. These virtual appliances can provide a quick look into specific vendor features and interfaces or can be deployed as a production solution that can utilize the resources available in your virtual infrastructure.

Hardware Load Balancing

When looking at enterprise-level deployments of web farms, you tend to see more hardware/appliance-based load balancers than anything else. The reason is that these are generally dedicated and clustered hardware devices whose sole purpose is to provide wire speed throughput and increased traffic performance. High-end hardware solutions from vendors such as F5, Cisco, and Citrix are built upon fast ASIC technology that provides considerable performance benefits over software solutions that must compete for CPU cycles.

Windows Network Load Balancing

When choosing your load-balancing solution, you will likely take many factors into consideration, among others the vendor, feature set, cost, ease of management, and compatibility with your virtual environment. With DNS round-robin and hardware load-balancing appliances, the solution is implemented outside the virtual environment. In most cases, this will require very little consideration to how the virtual infrastructure is designed and will focus more on the applications to be load-balanced, physical or virtual, and possibly the OS. However, if a software solution is chosen, there may be specific deployment methods you must follow, especially in a virtual environment.

Windows Network Load Balancing is one such example. NLB is installed as a feature on each host in what will soon be the web farm, pool, or cluster of servers. When configured, the cluster will present a single, virtual IP address on which each host will listen. Requests coming into the cluster's virtual IP will be directed to all the available cluster members and responded to by one of the members; all other cluster members will drop the request. Cluster members exchange heartbeat information to track the availability of other members. If a cluster member fails to respond (the default time is five seconds), they are tagged as unavailable, and requests are distributed to the remaining members.

Windows NLB has two modes of operation:

Unicast mode (default) In unicast mode, the Ethernet adapter's MAC address is changed to the MAC address of the cluster. That's correct—all cluster members have the same MAC address when in unicast mode. For this reason, each cluster member must have two Ethernet adapters to allow for intercluster communication.

Multicast mode When multicast mode is used, the cluster member's Ethernet adapters are again given the MAC address of the cluster; however, the original MAC address is retained. This allows intercluster communication without the requirement for a second Ethernet adapter. However, when multicast cluster members respond to ARP requests, they map their unicast IP address to the multicast MAC address; many routers do not support this and will require a static ARP entry for the requests to properly distribute among the nodes.

The NLB mode chosen will depend on factors such as the number of Ethernet adapters installed in cluster nodes, compatibility of network equipment, and communication requirements between cluster nodes. Because of requirements placed on networking equipment when running in multicast mode, the default mode of unicast has been the preferred method for deployment. In the following section, we will walk you through configuring Windows NLB in a vSphere environment.

Configuring vSphere for Windows Network Load Balancing

Windows NLB has been widely used because of its ease of installation, deployment, and configuration. However, if you have ever tried to install and configure NLB in a virtual infrastructure without doing a bit of research, you may have found that the load wasn't exactly balanced as you had expected. ESX hosts in a vSphere environment require additional configuration to allow NLB to function as designed.

When in unicast mode, requests to the NLB virtual IP address are flooded to all switch ports. Switched networks attempt to prevent this by having the switch learn the MAC addresses of hosts sending traffic through it. If allowed to learn the source port of the MAC address of your NLB virtual IP address, you could get stuck in a situation where all of your requests are sent to only one of your cluster nodes. To prevent the switch from learning the MAC address of your virtual IP address, NLB masks the MAC address on outbound traffic.

So, what's the problem? Well, you must remember that you are working at a couple of layers removed from the physical network. Your virtual machines must traverse the virtual network and hypervisor before reaching the physical network. This is where the problem occurs. ESX hosts attempt to help you out by sending reverse ARP (RARP) packets to ensure your VMs can be located quickly. This is important, especially when you quickly move a VM from one ESX host to another such as during a vMotion or VMware HA recovery. However, this means that the masking NLB provides is sabotaged by ESX.

The good news is you can work around this issue by making a few minor configuration changes in ESX. During the Source Discovery proof of concept, we will walk you through the steps taken by the web administrators to configure Windows Network Load Balancing. We will also outline the steps taken by the Source Discovery ESX administrators so you can be sure you are correctly deploying WNLB services on top of vSphere.

Windows Network Load Balancing is a great way to add redundancy to an IIS environment. Web servers must be able to respond quickly to user requests; otherwise, visitors might be turned away because of slow-loading pages. By deploying multiple web servers in a scaled-out fashion and load balancing the requests among the web servers, you are able to maximize the efficiency of your hardware and virtual machines.

With the potential for deploying many virtual machine instances very quickly, it is easy to lose sight of the critical management tasks such as the security of the operating systems and applications. Web servers in general must be given an extra once-over to ensure they don't become some malicious attacker's jump box. Let's take a look at some of the key points to securing IIS.

Securing Virtualized Internet Information Services

For most of us, the world, as we're accustomed to it, is secure. Intranets are secure and friendly, and you know what's out there, right? Sure, but what about those externally facing websites? Sooner or later most of us will have to administer those systems, but how do you ensure you secure them properly? It wasn't too long ago that we considered IIS servers on the Internet to be taboo. Is this still the case? Let's find out.

Versions of IIS prior to IIS version 6.0 came packed with features and capabilities that were supposed to make it easy to use. What really happened was that it quickly became a target for malicious attacks. Why not, right? After all, you had such a large attack surface. IIS 6.0 (available with Windows Server 2003) changed that in two ways. First, it was not installed by default, as it was in Windows 2000. Second, IIS was installed with only the ability to serve static content. IIS 7 has taken security to a level where it can now fully compete with web servers such as Apache.

IIS 7 has been rewritten with security as a major design improvement. The new modular design means that the default installation is a minimalistic approach to deploying a web server. Installing IIS does not mean you will get any bells or whistles. If you want fancy dynamic web pages, ASP.NET, or CGI support, you will need to do a bit more work, but you can serve static pages, monitor the server traffic, and provide compression for static content. So, what do you get from this new approach?

- The fewer modules you install to serve your web content, the smaller your attack surface is.
- Having fewer components installed on your IIS server means you have less to manage and maintain.
- Fewer components running on your web server means less memory utilization, which equates to better performance.

Besides going to a fully modular design, other security enhancements include built-in request filtering. Request filtering processes requests on the fly and can filter based on filename extensions, HTTP verbs, size, namespace, and sequences. If you have used the add-on for previous

versions called URLScan, you are familiar with request filtering, but now it's built in to the application. In the following sections, we'll talk about adding security to your web service environment.

Authentication

All of the previously available authentication schemes are still available in IIS7:

AD client certificate Active Directory client certificates can be mapped to users for authentication, allowing them to be automatically authenticated without using any other schemes.

Anonymous Anonymous allows any user to access the public content on your site without providing a username or password. This is the default enabled authentication scheme in IIS 7.

Basic Users provide a username and password to access content. Basic authentication sends user passwords across the network using a weak level of encryption. Basic authentication is usually augmented by ensuring the connection is protected via a Secure Sockets Layer (SSL) connection.

Digest Digest uses Windows domain controllers to authenticate users requesting access to content; it is usually deployed when improved security over basic authentication is required. Browsers must support the HTTP 1.1 protocol to use digest authentication.

Forms Unauthenticated users are redirected to an HTML form where they can enter their username and password. Once authenticated, users are redirected to the originally requested page. As in basic authentication, it is recommended that SSL be used to encrypt the communications between the client and web server because forms authentication transmits the username and password in clear text.

Windows NTLM or Kerberos security providers are used to authenticate clients when Windows authentication is used. Windows authentication's default setting is to negotiate the security provider or have the client select which provider to use. Some environments will want to enforce one or the other. This scheme is best suited for intranet environments.

On its own IIS can be made quite secure. That doesn't mean you are going to simply place your IIS server on to the Internet and hope things go well. Chances are they won't, even with the great strides that have been made to secure the IIS application, the Windows OS, and the work you do to harden your build. Security engineers (good and bad) continue to find exploitable holes. This is evident in Microsoft's "Patch Tuesday." Regardless, you have to find a way to securely publish your sites to the outside world. This is where perimeter security comes into play.

Using VMware's vShield Zones Firewall

Hardware firewalls such as those from vendors such as Cisco and Checkpoint have traditionally protected intranets from the world beyond. Firewalls are usually dedicated devices that sit outside of the compute environment and permit or deny network traffic in or out of an environment based on a set of rules. In its simplest form, a firewall can sit in front of your IIS server and only allow ports 80 and 443. This prevents unauthorized users from accessing other open ports such as 135 (MS-RPC), 445 (SMB), or 3389 (RDP), as shown in Figure 5.16.

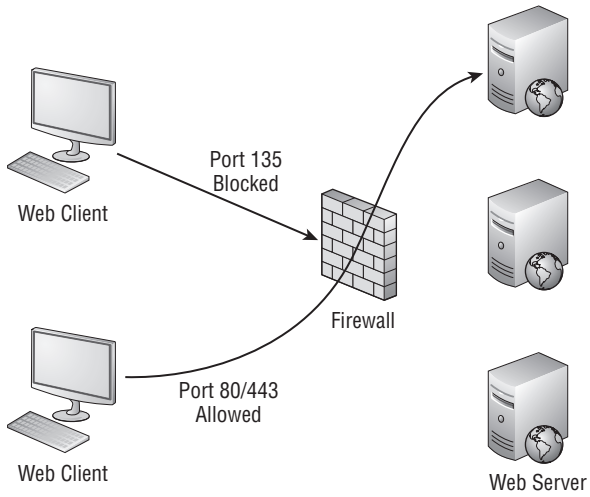

FIGURE 5.16
Firewalls permit or deny traffic on specific ports to hosts on the DMZ.

Software firewalls work on the same principles as hardware firewalls but use existing computing resources as opposed to dedicated hardware devices. Software firewalls can be installed as an additional application on an existing OS or in the form of a virtual appliance. Microsoft Internet Security and Acceleration Server is a good example of a software-based firewall solution that provides secure access to internal resources and allows for application-specific rules, URL filtering, and intrusion prevention. Microsoft has deployed the newest version of this application, now named Microsoft Forefront Threat Management Gateway (TMG).

vShield Zones

VMware has not been left out of this space. The acquisition of Blue Lane Technologies by VMware in 2008 opened the door for virtual infrastructure protection using vShield zones. With VMware vShield, you now have the ability to provide integrated protection for virtualized data centers. Deploying vShield zones gives you an inventory of the operating systems, applications, and open ports within your virtual data center. You can then use this information to segregate traffic flow into logical trust zones and set up security policies without having to fork the traffic off through an external device. This is especially useful when managing a virtual IIS infrastructure that must be externally available. You no longer need to deploy a vSphere cluster or host for DMZ-only access. Using vSphere zones, you can bring the DMZ into the virtual environment and fully utilize your computing resources throughout the entire vSphere environment. Using vShield zones you are able to segregate the information as shown in Figure 5.17.

The components that make up vShield zones environment are as follows:

vShield A virtual appliance located on each vSphere host used to inspect traffic flow and provide firewall protection

vShield Manager Manages all the distributed vShield instances by providing monitoring, centralized configuration, and software updates

Once deployed, vShield sits in between your protected virtual machines and the external network interfaces. This allows vShield to intercept each network packet and enforce the policies that have been created for the particular security zone. For more information on how vShield zones can help safeguard your virtual infrastructure, see the VMware website: `www.vmware.com/products/vshield-zones`.

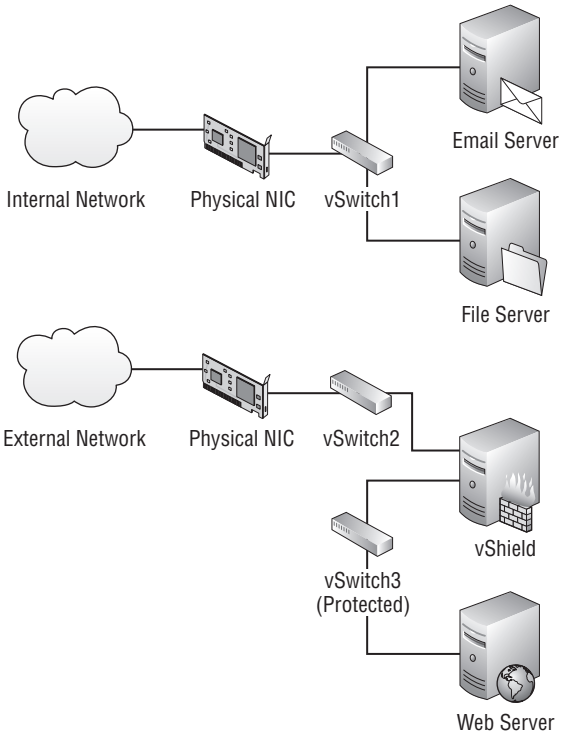

FIGURE 5.17
Protected virtual machines are isolated on the virtual switch with no direct access to the external network.

Virtualizing IIS Proof of Concept

As Source Discovery continues virtualizing its tier-1 applications, it sees the various IIS servers it currently has running in production as low-hanging fruit. The company has decided that these miscellaneous web servers are quick wins to prove how virtualization can help them achieve quick gains in its server consolidation project. The IIS servers scoped for this round of migrations are front ends to many critical single-tier web applications. The requirements include a load-balancing solution, monitoring, and virtual machine–level backup and recovery because of the highly customized configurations involved in building these systems.

The Source Discovery virtual infrastructure team will deploy virtual machine templates with the virtual hardware as required by the web administration team. The web administrators will then install IIS and enable the remote management of IIS, which is the default administration method practiced at Source Discovery. The web administration team will configure IIS

using the optimization tips outlined earlier in the chapter. To provide a load-balanced environment, the web team will add the Windows Network Load Balancing feature and work with the virtual infrastructure team to make sure ESX is properly configured to allow WNLB to work as designed. While the web administration team finalizes the web application configuration, the virtual infrastructure team will install and configure VMware AppSpeed to provide SLA monitoring and trending.

With template creation and deployment covered in Chapter 3, you can pick up at the point of installing IIS and working your way through the remaining requirements.

Installing IIS

The following procedure shows how to install IIS 7 with the default modules:

1. Launch the Server Manager, located on the Start menu under Administrative Tools (Figure 5.18).

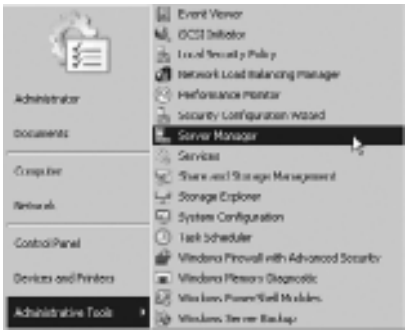

FIGURE 5.18
Server Manager is located under the Administrative Tools submenu.

2. On the Server Manager dashboard or by selecting the Roles object from the left pane, click the Add Roles link to start the Add Roles Wizard (Figure 5.19).

FIGURE 5.19
Server Manager roles management

3. The first screen you see in the wizard asks you to verify a few of the points we made earlier; click Next to continue (Figure 5.20).
4. On the Select Server Roles screen, select Web Server (IIS), and click Next (Figure 5.21).

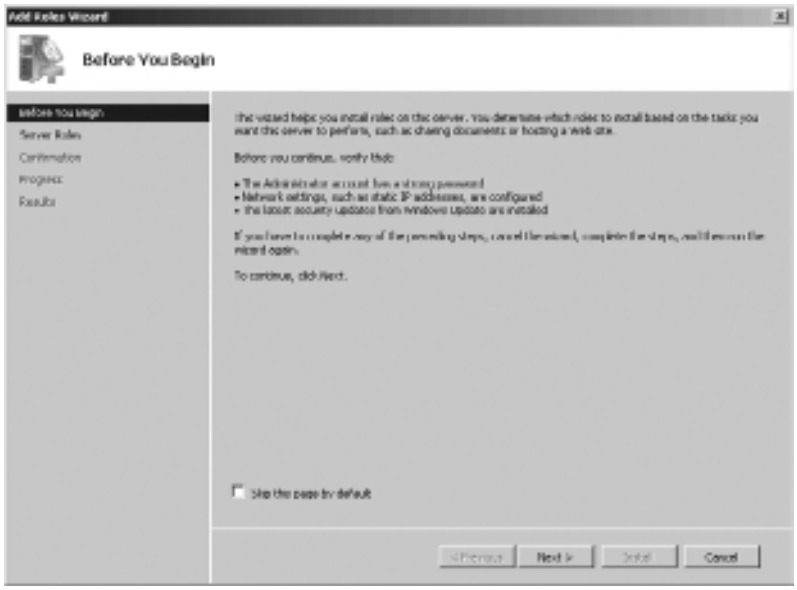

FIGURE 5.20
Verify that the administrator account, IP settings, and security updates have been reviewed before continuing.

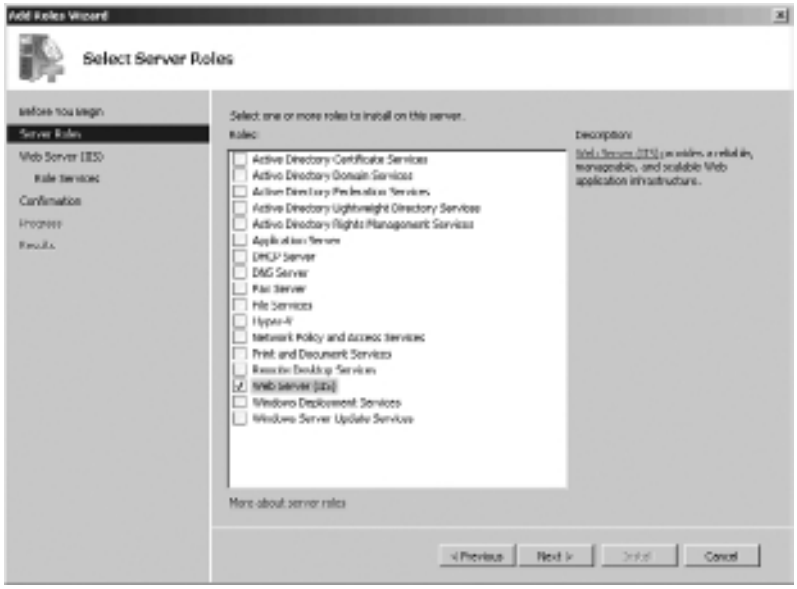

FIGURE 5.21
Select Web Server screen

5. The Web Server (IIS) screen provides a few links with more information regarding the IIS web server; you can review this now or return to it later for reference. Click Next to continue (Figure 5.22).

FIGURE 5.22 Introduction to web server and links for additional information

6. The Select Role Services screen lets you review the default modules installed with IIS as well as install any additional modules that are available out of the box. Highlighting modules will give a brief description in the upper-right corner. Leave the default selections selected, and click Next (Figure 5.23).

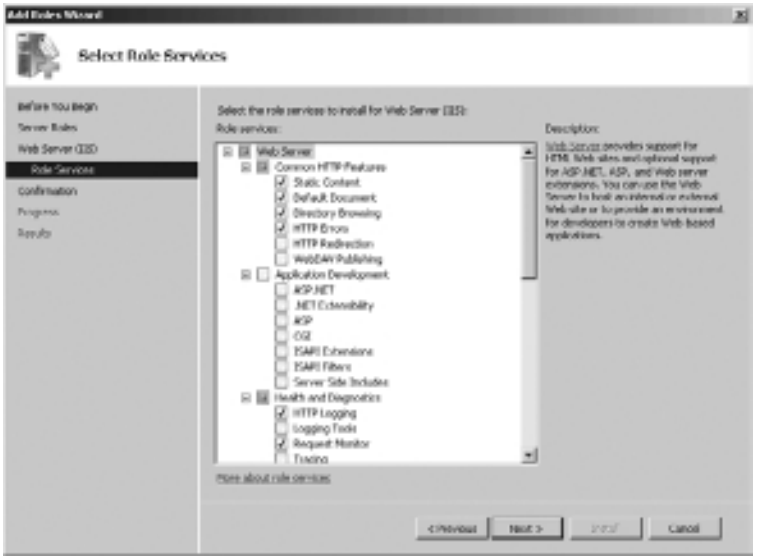

FIGURE 5.23 Minimal role services selected for IIS

7. Review the settings you have chosen, and click Install (Figure 5.24).

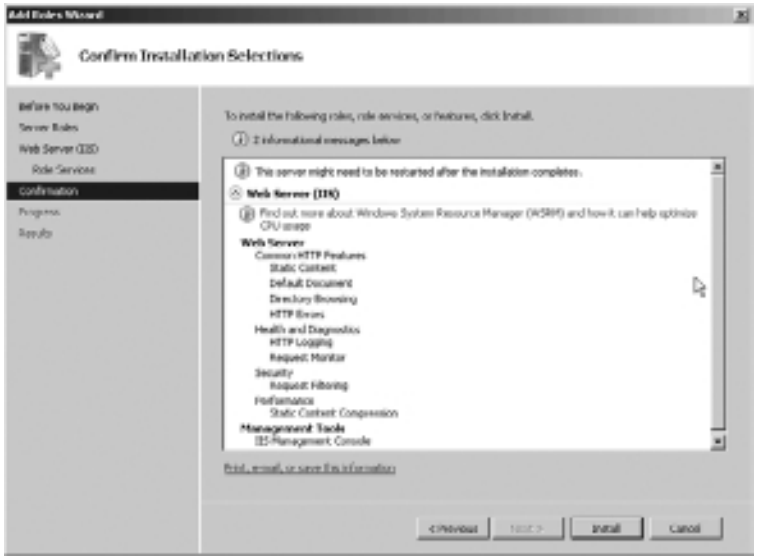

FIGURE 5.24
Review the chosen role components to complete the installation.

8. Once the installation is complete, the results screen gives a summary of the changes made to the system; click Close to complete the process (Figure 5.25).

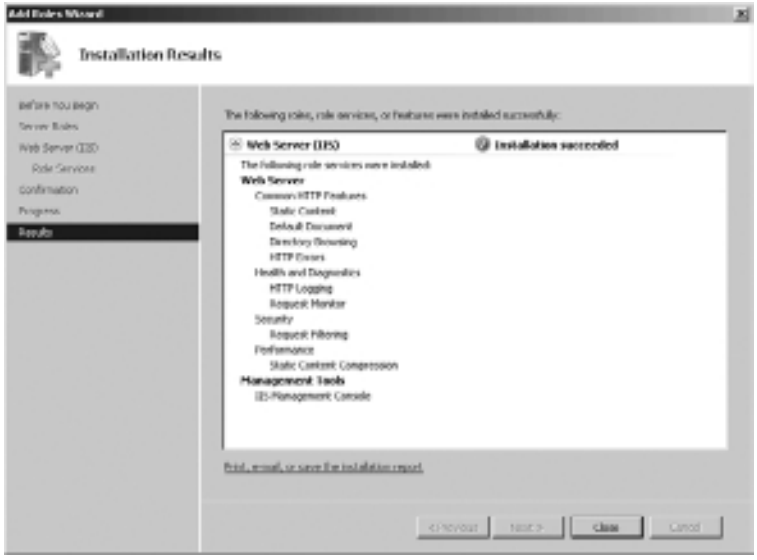

FIGURE 5.25
The results screen shows a successful installation.

9. To verify the installation, open a browser, and navigate to `http://localhost`. The default IIS 7 page confirms IIS has been installed and is functioning (Figure 5.26).

FIGURE 5.26
Default landing page for IIS 7 installation

10. To review the default settings, open the Internet Information Services Manager from the Administrative Tools submenu on the Start menu (Figure 5.27). Figure 5.28 shows the GUI of IIS Manager.

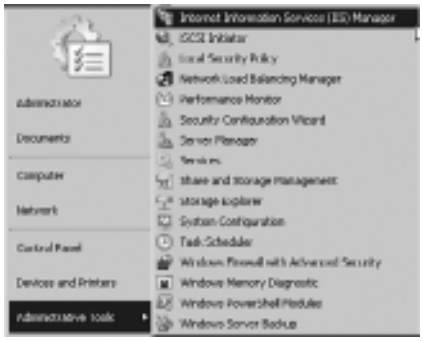

FIGURE 5.27
Launch IIS Manager from the Administrative Tools submenu.

With IIS installed, tested, and confirmed working, the web team moves on to enabling the remote administration of IIS. If you recall, the step-by-step process for enabling remote administration when deploying in the Server Core version of Windows Server 2008 was described earlier in the chapter. The following process describes enabling remote IIS management in the full version of Windows Server 2008.

FIGURE 5.28
IIS Manager provides an extensible GUI for managing all aspects of IIS.

Remote Management in a Full Installation

Follow this procedure to enable remote management in a full installation:

1. Install the IIS Remote Management service using the Add Role Services Wizard (Figure 5.29).

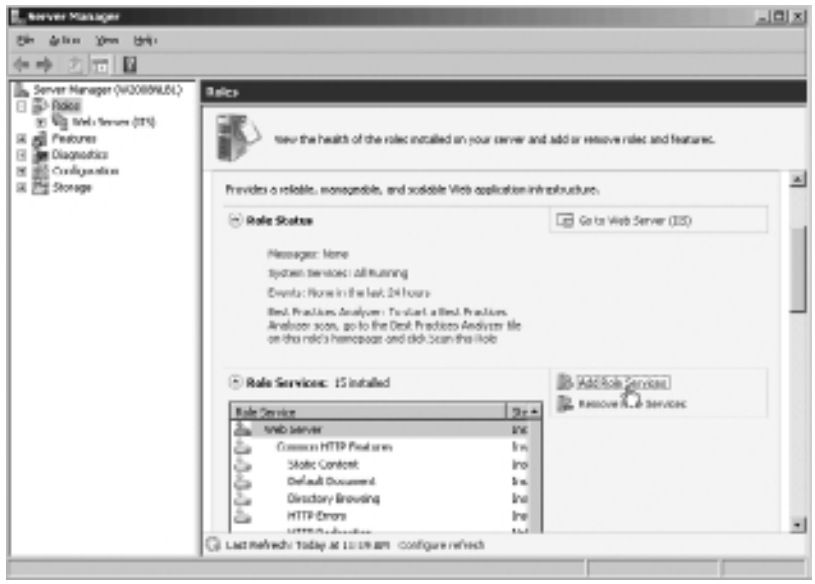

FIGURE 5.29
Launch the Add Role Services Wizard.

2. Select Management Service from the list of available role services, and click Next (Figure 5.30).

FIGURE 5.30
Select the Management Service role.

3. Confirm the role services to be installed, and click Install.
4. Open IIS Manager, and select the local server.
5. Select Management Service from the features pane under the Management section (Figure 5.31).

FIGURE 5.31
Open the Management Service feature in IIS Manager.

6. Select the Enable Remote Connections box, and click Apply in the Actions pane (Figure 5.32).

7. Once enabled, the service can be started by clicking the Start link from the Actions pane (Figure 5.32).

FIGURE 5.32
Enable remote management, and start, stop, or restart the service.

8. Set the Web Management Service to start automatically when the server boots up using the following command (the space between the = and `auto` is required):

```
sc config WMSVC start= auto
```

Connecting to the Remote Server

Use the following steps to connect to the remote server:

1. Open IIS Manager, click the connection pull-down menu on the Connections pane, and click Connect To A Server (Figure 5.33).

FIGURE 5.33
Connections menu

2. Enter the name of the remote server or website to connect to, and click Next (Figure 5.34).

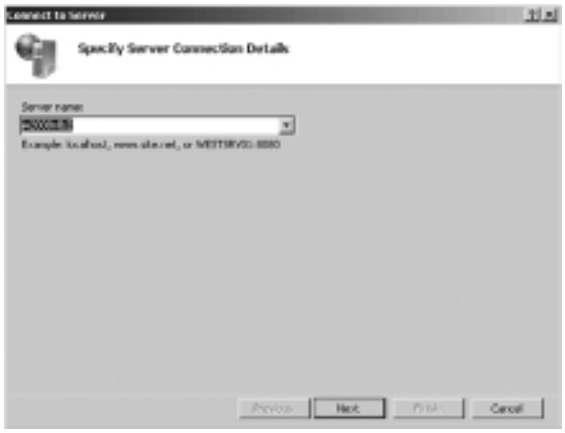

FIGURE 5.34
Enter the name of the remote system to connect to.

3. Provide credentials to connect to the remote system, and click Next (Figure 5.35).

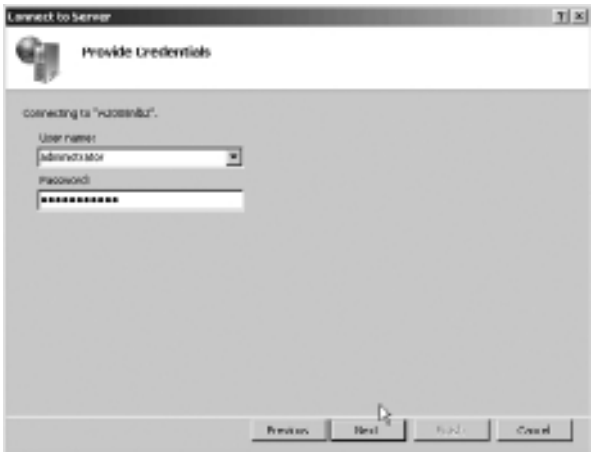

FIGURE 5.35
Provide credentials to access the remote system.

4. If desired, change the connection name, and click Finish to complete the connection (Figure 5.36).

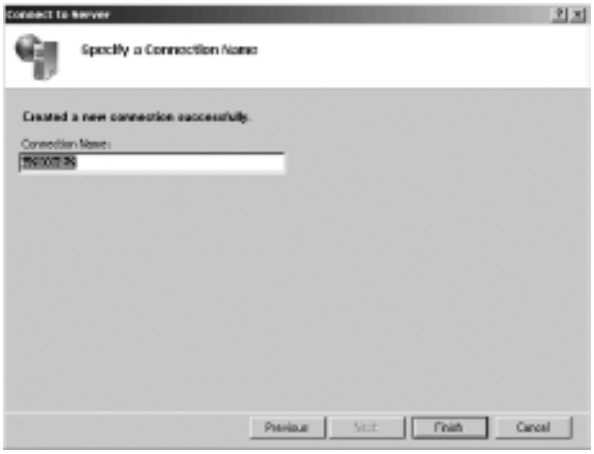

FIGURE 5.36
Successful connection created

5. The new connection will appear alongside the local connections (Figure 5.37).

FIGURE 5.37
Local and remote connections appear in a single IIS Manager window.

Configuring ESX for Windows Network Load Balancing

With IIS installed and accessible remotely, the Source Discovery web administration team can begin the setup process for Windows Network Load Balancing. Following this process, the web team has built a two-node WNLB cluster:

1. Windows NLB will require an extra virtual Ethernet adapter for NLB use. If this wasn't originally called out in the requirements during deployment, the additional NIC can be added using the vCenter client. Open the virtual machine properties by right-clicking the virtual machine and selecting Edit Settings (Figure 5.38).

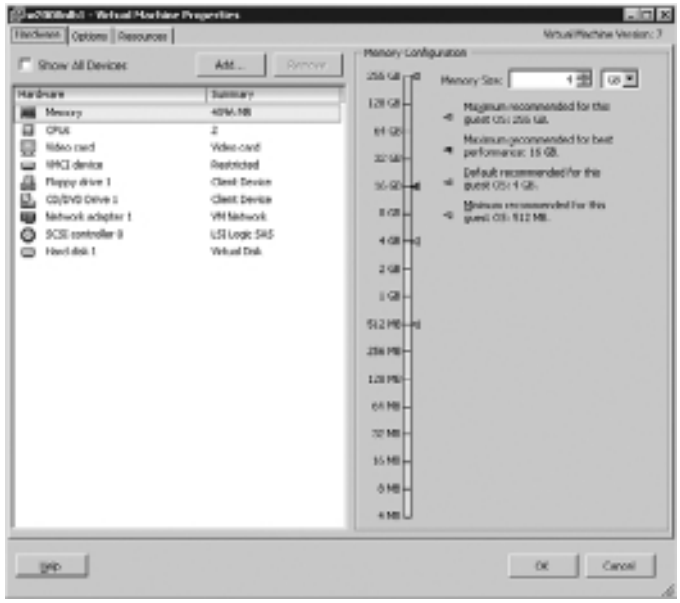

FIGURE 5.38
Virtual Machine Properties window

2. Click Add to launch the Add Hardware Wizard.
3. Select Ethernet Adapter, and click Next (Figure 5.39).

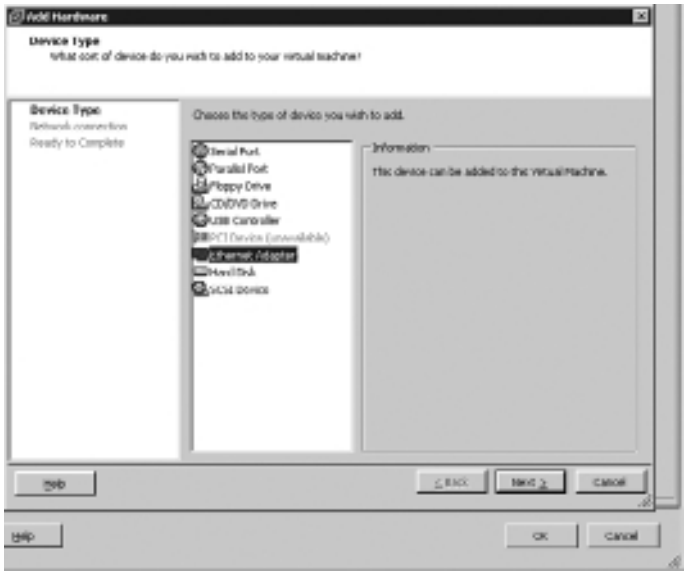

FIGURE 5.39
Add a second virtual Ethernet adapter.

4. Select the type of adapter; we recommend VMXNET 3. Select the network that will connect the VM to the appropriate network, and click Next (Figure 5.40).

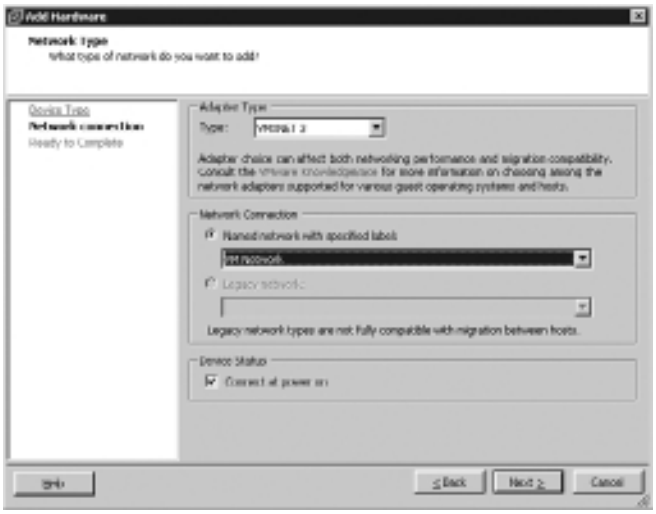
Add Hardware
Network Type
What type of network do you want to add?
Device Type
Network connection
Ready to Complete
Adapter Type
Type:
Adapter choice can affect both networking performance and migration compatibility. Consult the VMware KnowledgeBase for more information on choosing among the network adapters supported for various guest operating systems and hosts.
Network Connection
Named network with specified label:
Legacy network:
Legacy network types are not fully compatible with migration between hosts.
Device Status
Connect at power on
Help
< Back
Next >
Cancel

FIGURE 5.40
Select the Ethernet adapter type and network.

5. Review the summary screen, and click Finish (Figure 5.41), once added the virtual hardware dialog box will show the newly added Ethernet adapter in bold (Figure 5.42).
6. Power on the virtual machines, and log on with an administrator account.

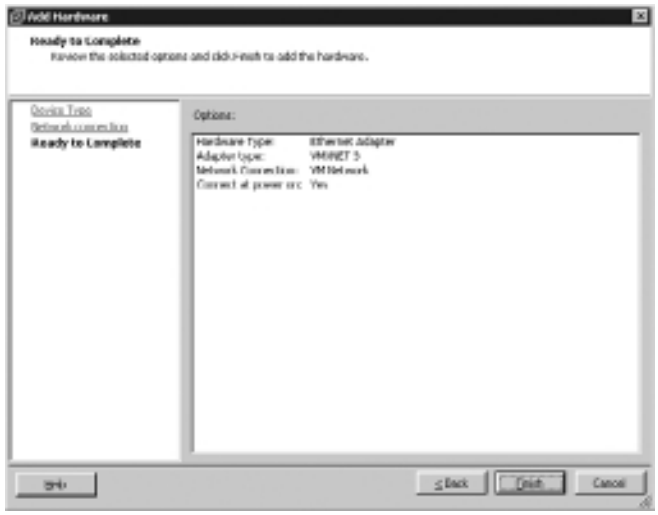

FIGURE 5.41
Hardware wizard summary screen

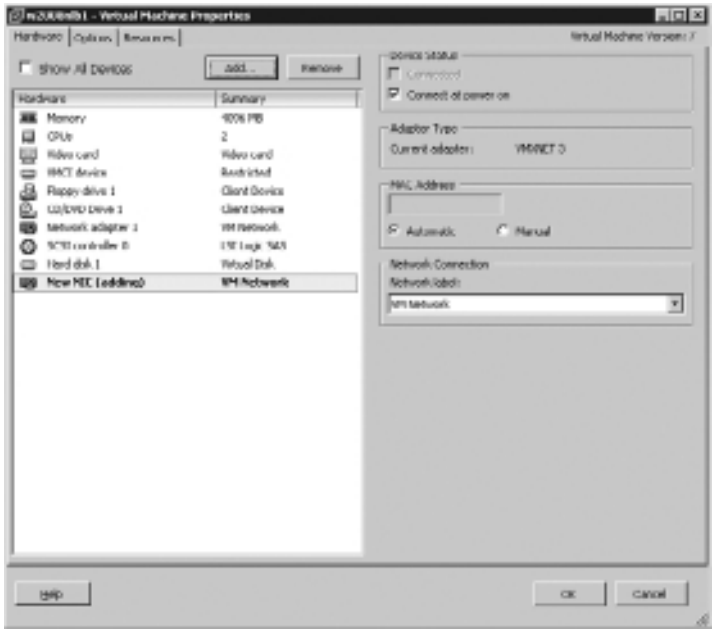

FIGURE 5.42
Virtual Ethernet adapter ready to be installed in the virtual machine

7. Open the Control Panel from the Start menu, and select the Network And Internet category (Figure 5.43).

FIGURE 5.43
Network And Internet control panel

8. Open the Network And Sharing Center, and click the Change Adapter Settings link (Figure 5.44).

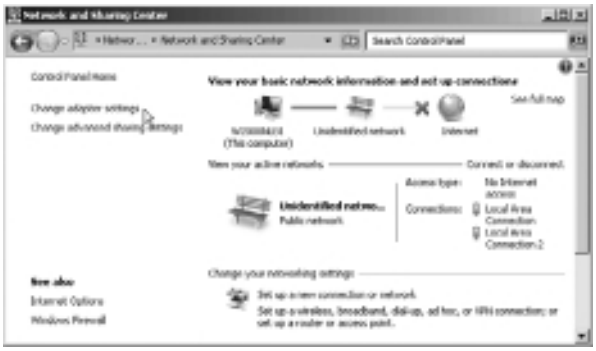
Network and Sharing Center
Network ... ▸ Network and Sharing Center ▸
Search Control Panel
Control Panel Home
Change adapter settings
Change advanced sharing settings
View your basic network information and set up connections
See full map
(This computer)
Unidentified network
Internet
View your active networks
Connect or disconnect
Access type:
No Internet access
Connections:
Local Area Connection
Local Area Connection 2
Public network
Change your networking settings
Set up a new connection or network
Set up a wireless, broadband, dial-up, ad hoc, or VPN connection; or set up a router or access point.
See also
Internet Options
Windows Firewall

FIGURE 5.44
Network And Sharing Center

9. Identify which Ethernet interface will be used for NLB and management, and label it using a descriptive name (Figure 5.45).

FIGURE 5.45
Give each interface a descriptive name for easy identification.

10. Configure the static IP settings for each network interface; you can verify connectivity to each interface by using `ping.exe` or `tracert.exe` (Figure 5.46).

FIGURE 5.46
Verify each interface is accessible over the network.

11. Once the networking hardware has been configured, you can begin configuring Network Load Balancing. Launch Server Manager from the Administrative Tools submenu on the Start menu (Figure 5.47).

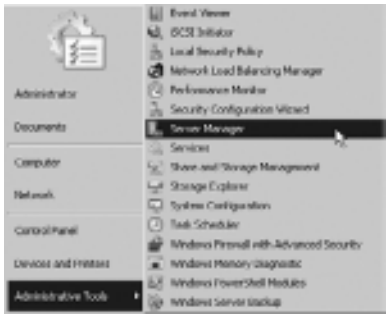

FIGURE 5.47
Open Server Manager from the Administrative Tools submenu.

12. From the left side menu in Server Manager, select Features.
13. Click Add Features to kick off the Add Features Wizard (Figure 5.48).

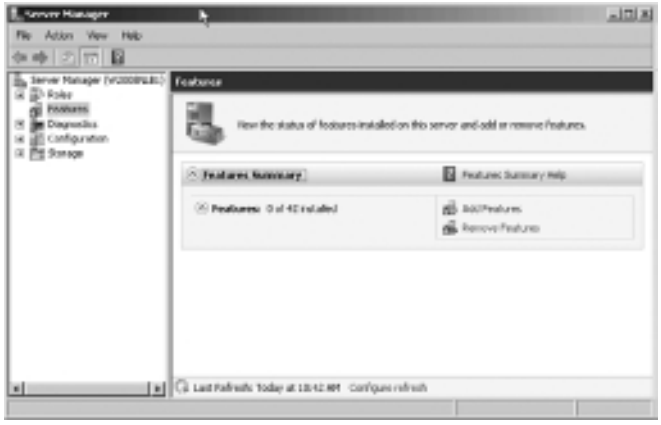

FIGURE 5.48
Use Server Manager to install server features.

14. In the Features selection field, locate and select the Network Load Balancing check box (Figure 5.49).

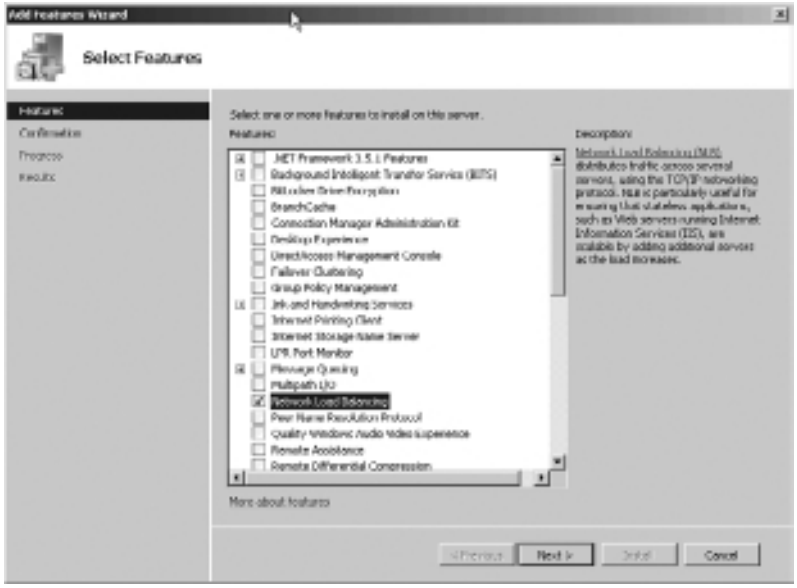

FIGURE 5.49 Select Network Load Balancing from the feature selection pane.

15. The confirmation field provides any informational messages, such as letting you know that a restart may be required. Click Install to continue (Figure 5.50).

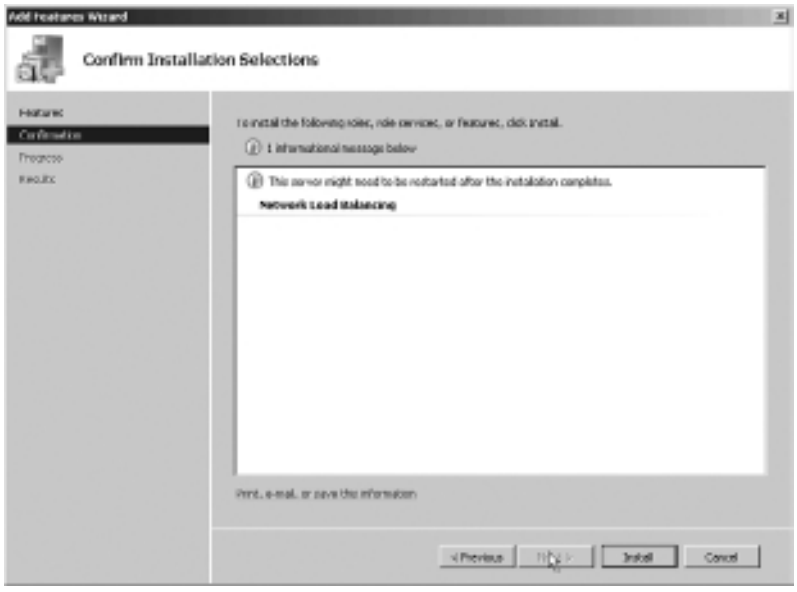

FIGURE 5.50 Confirm the features to be installed.

16. Once complete, click Close to complete the wizard (Figure 5.51).

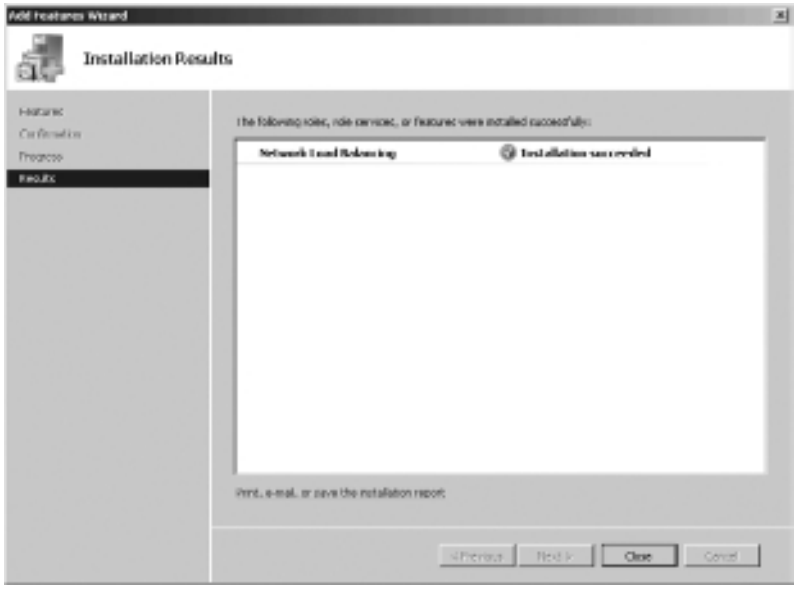

FIGURE 5.51 Review the summary page to complete installation.

17. By adding the NLB feature, we've also added the management tools for NLB. Open the Network Load Balancing Manager from within the Administrative Tools submenu under the Start menu (Figure 5.52).

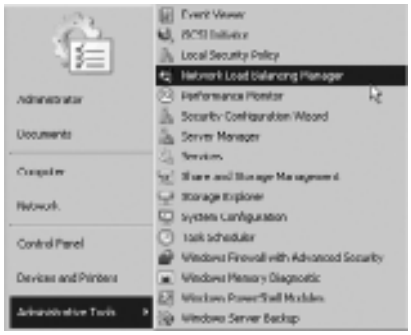

FIGURE 5.52 Network Load Balancing Manager located in the Administrative Tools submenu

18. Network Load Balancing Manager will load with no connections. You could remotely connect to and manage or configure an existing NLB cluster, or you can create a new cluster. Because this is your first deployment of NLB, right-click the Network Load Balancing Clusters node in the left pane, and select New Cluster (Figure 5.53).

FIGURE 5.53
Create a new load-balancing cluster.

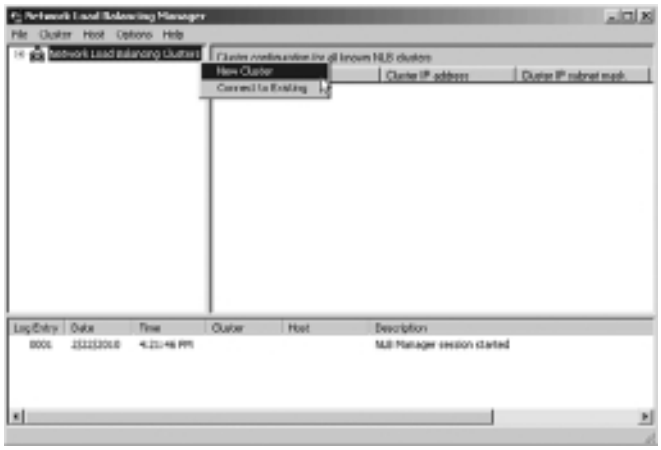

19. Enter the name or IP address of the first host to be added to the cluster, and click Connect (Figure 5.54); you must have administrative rights to the machine.

FIGURE 5.54
Enter the host-name or IP address of the first member in the NLB cluster.

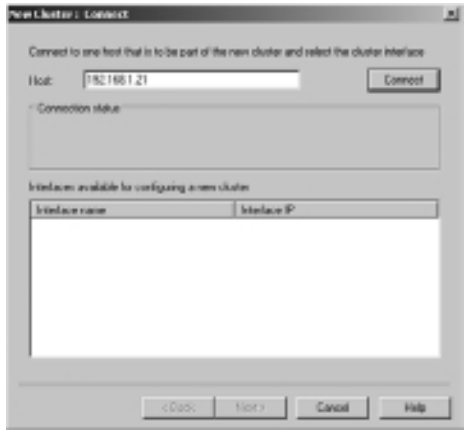

20. Once connected, the interfaces will appear in the interfaces pane at the bottom of the window; select the NLB interface, and click Next (Figure 5.55).

21. The default settings may be used.

    The Priority (Unique Host Identifier) value specifies a unique ID for each host in the cluster. Lower-priority hosts will handle cluster network traffic not covered by a port rule.

    The Initial Host State setting determines whether NLB will be started or stopped after a restart of the OS and whether an OS restart will reenable the NLB service if you had previously had it suspended.

Make your selections, and click Next (Figure 5.56).

FIGURE 5.55
Connect to the first host in the cluster.

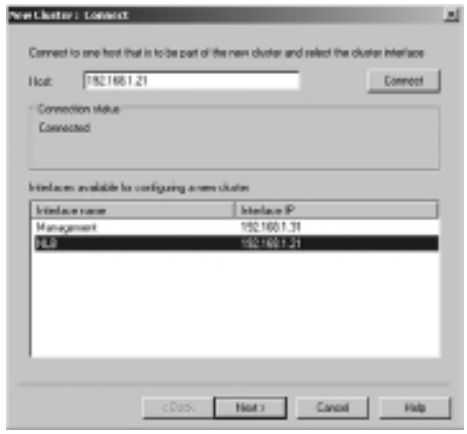

FIGURE 5.56
Set the host parameters for node 1.

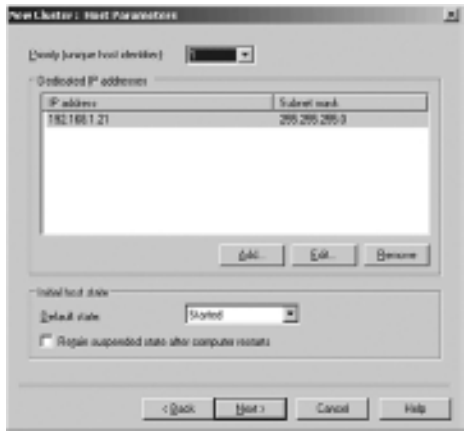

22. In the Cluster IP Addresses dialog box, click Add, and enter the IP address where all communication to the NLB cluster will be directed. This IP address is shared by all members of the load-balancing cluster. Click Next to continue (Figure 5.57).

FIGURE 5.57
Set the IP address for the cluster.

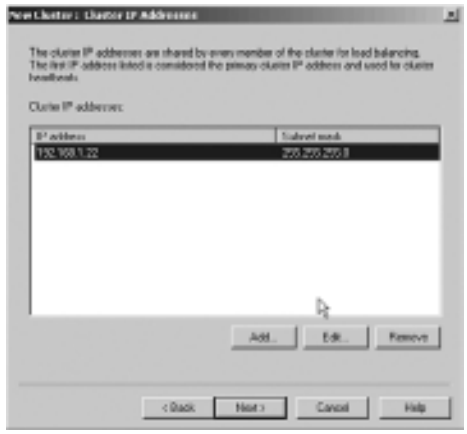

23. The Cluster Parameters dialog box allows you to choose the Cluster IP if multiple IPs were configured and provide a fully qualified name for the NLB cluster. The Cluster Operations mode gives you the option to select unicast or multicast mode. To provide the highest level of compatibility, you will leave the default selection of Unicast and click Next (Figure 5.58).

FIGURE 5.58
Configure the cluster parameters.

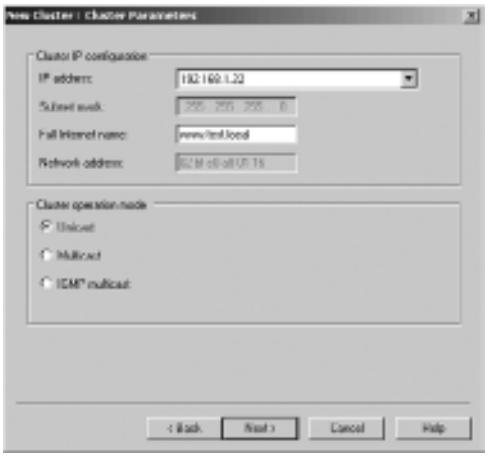

24. You can review the settings on the Port Rules dialog box, but because you will only be configuring a two-node cluster and you want traffic evenly distributed to both nodes, leave the default rule, and click Finish (Figure 5.59).

FIGURE 5.59
Cluster port rules allow you to direct requests on specific ports to specific cluster nodes; the default spreads the load to all nodes.

Once completed, the cluster and any attached hosts will appear in the left pane and appear as green (started), gray (suspended), or red (stopped), as shown in Figure 5.60.

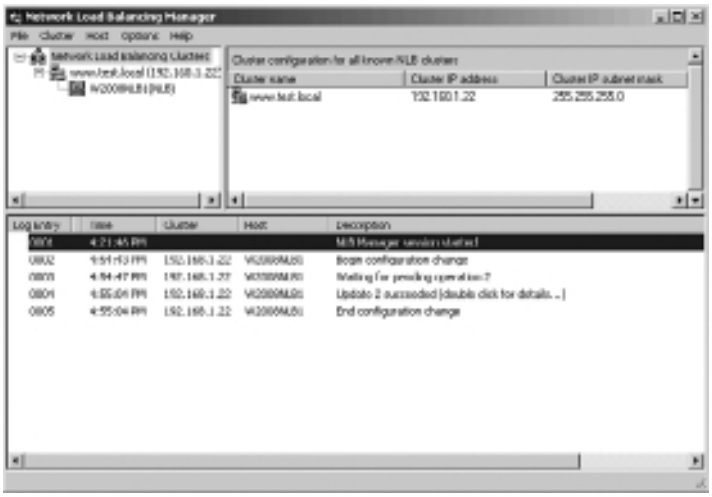

FIGURE 5.60
Cluster configured with a single node online

25. To add the second node, right-click the cluster, and select Add Host To Cluster (Figure 5.61).

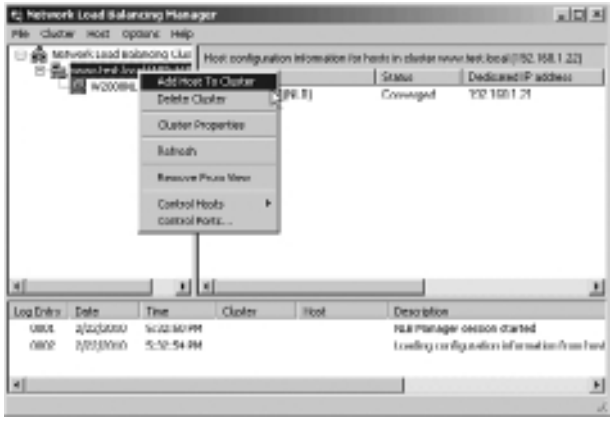

FIGURE 5.61
Add Host To Cluster option

26. As when creating the cluster, enter the host name or IP address, and click Connect (Figure 5.62).
27. Select the interface to use for NLB traffic, and click Next.

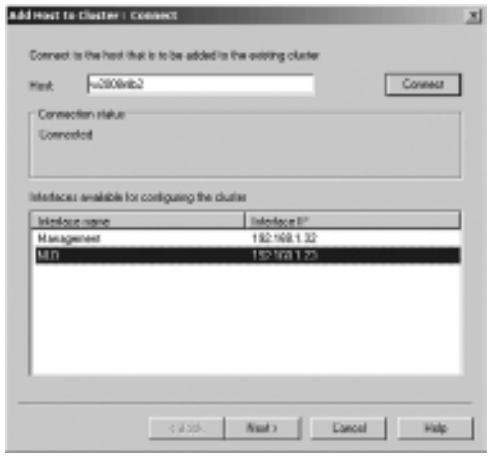

FIGURE 5.62
Connect to the second cluster node.

28. The next few screens will be almost identical to those during NLB cluster creation. Continue with the defaults, and click Finish on the Port Rules screen (Figure 5.63).

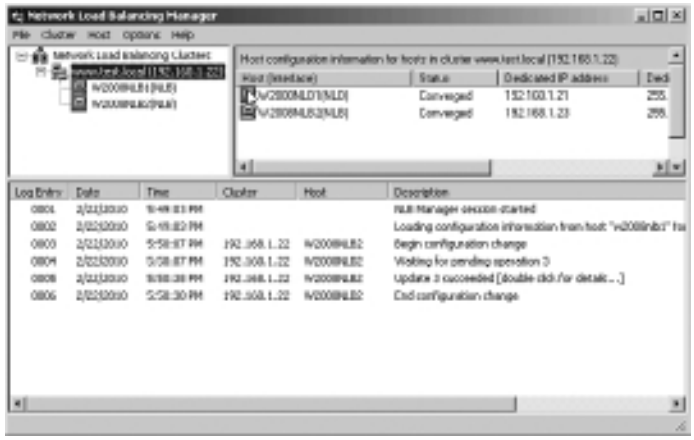

FIGURE 5.63
Two-node NLB cluster

CONFIGURING THE ESX HOSTS FOR WINDOWS NETWORK LOAD BALANCING

With the NLB cluster complete, the ESX administration team can make the changes to the underlying ESX host(s) to ensure that load balancing is occurring correctly between multiple hosts:

1. Log into the VMware vSphere Client, and locate and select the ESX host(s) that house the NLB VMs.
2. Click the Configuration tab, and select Networking (Figure 5.64).

FIGURE 5.64
Network settings for vSphere host

3. Locate the virtual switch that the NLB VMs cluster interface is connected to, and select Properties.
4. On the Ports tab of the virtual switch properties, select vSwitch, and click Edit (Figure 5.65).

FIGURE 5.65
Cluster interface vSwitch properties

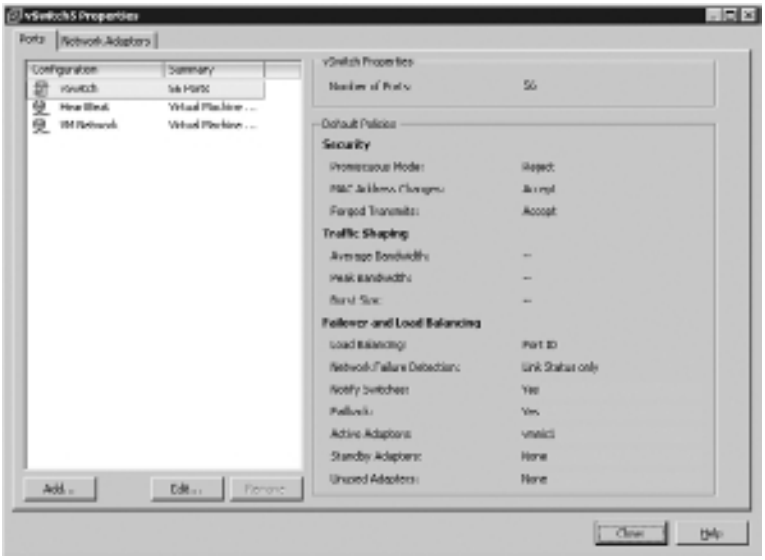

TIP To prevent RARP packet transmission at the port group level as opposed to the switch level, select the port group on step 4, and click Edit. The rest of the process is the same as at the switch level.

5. Click the NIC Teaming tab, and set the Notify Switches policy exception to No (Figure 5.66).
6. Click OK to close the properties dialog boxes.

FIGURE 5.66
The Notify Switches setting determines whether RARP packets are sent by the vSphere host

Installing VMware AppSpeed

The Source Discovery ESX administration team has begun expanding the vSphere tools they use in their production environments. They have recently begun looking into the VMware AppSpeed tool to show the in-depth reporting capabilities that the vSphere infrastructure provides. The web team will be the first group to begin taking advantage of this tool and have requested that their new VMs be monitored using AppSpeed. The ESX administrators followed this procedure to deploy the AppSpeed collectors and probes on the required ESX hosts:

1. Once AppSpeed is downloaded and extracted to a workstation running the vCenter client, launch the vCenter client and log into your vCenter Server. From within the vCenter client, select File ➢ Deploy OVF Template (Figure 5.67) to begin the wizard.

FIGURE 5.67
Launch the Deploy OVF Template Wizard.

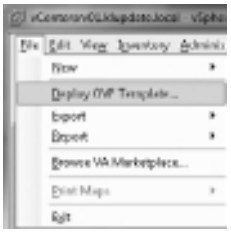

2. On the Source screen, select the radio button Deploy From File, and click the Browse button to locate the OVF file you have downloaded and unzipped, as shown in Figure 5.68; then click Next.

FIGURE 5.68
Deploy OVF Template Wizard's Source screen

Deploy OVF Template
Source
Select the source location.
Source
OVF Template Details
Name and Location
Ready to Complete
Deploy from file:
Browse...
Choose this option if the source OVF template (*.ovf) is on the local file system.
For example, your C: drive, a network share, or a CD/DVD drive.
Deploy from URL:
Choose this option to download the OVF template from the Internet and enter a
URL such as http://www.example.com/template.ovf
Help
< Back
Next >
Cancel

3. The OVF Template Details screen (Figure 5.69) will display information read from the OVF file and display them for you to understand the requirements. Review the requirements, and click Next to continue.

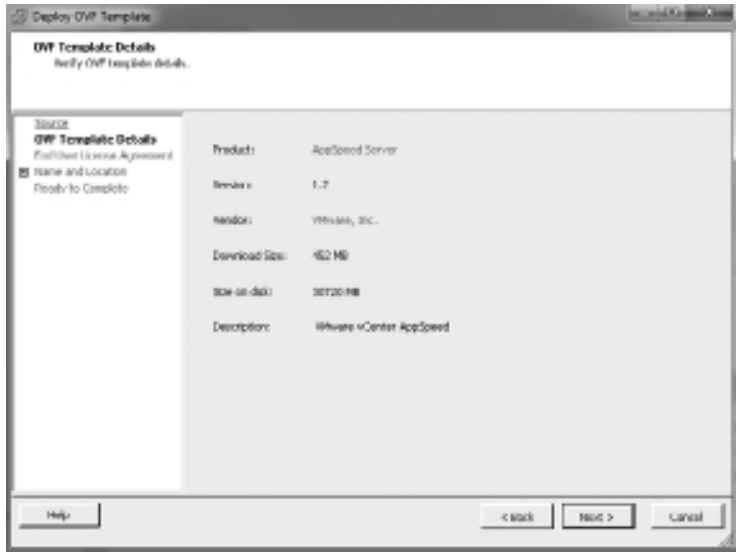

FIGURE 5.69
OVF template details

4. After reading the end user license agreement, displayed in Figure 5.70, click Accept and Next to continue.

FIGURE 5.70
End user license agreement

5. Give the AppSpeed Server a name (AppSpeed Server is fine), and select the location within your data center to place the virtual machine, as shown in Figure 5.71.

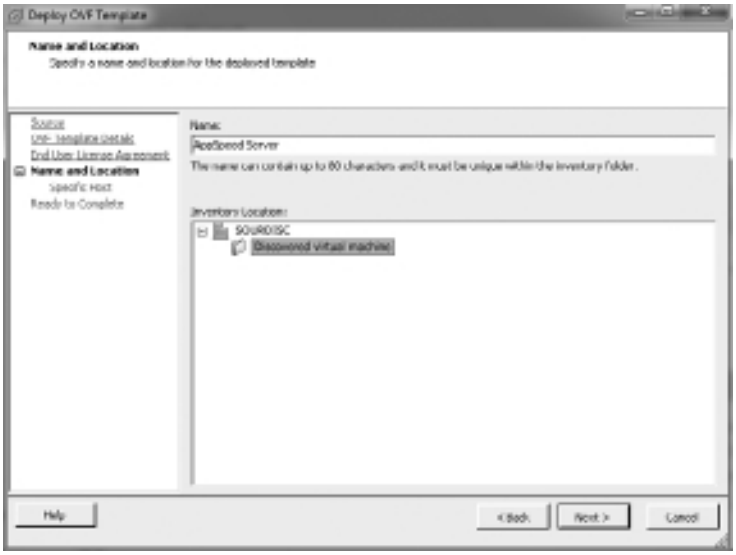

FIGURE 5.71 Specify the name and location for the AppSpeed appliance.

6. If DRS is enabled and set to manual in your vSphere cluster, you must assign the virtual machine to a host; in Figure 5.72, the host europa.klupdate.local has been selected. If DRS is set to automatic or partially automatic, the host will be automatically assigned.

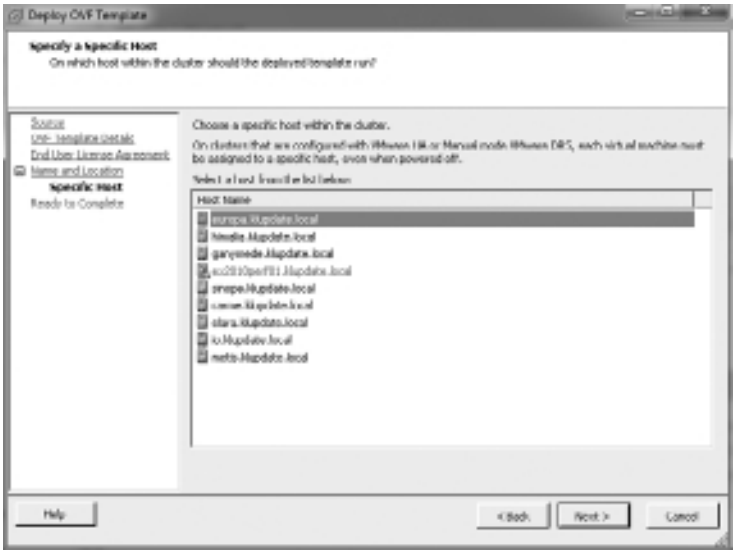

FIGURE 5.72 Specify a host for the virtual machine.

7. On the Datastore screen (Figure 5.73), select a location with plenty of disk space, preferably one that has access from multiple hosts to allow for vMotion, DRS, and HA.

FIGURE 5.73
Select a datastore for the virtual machine.

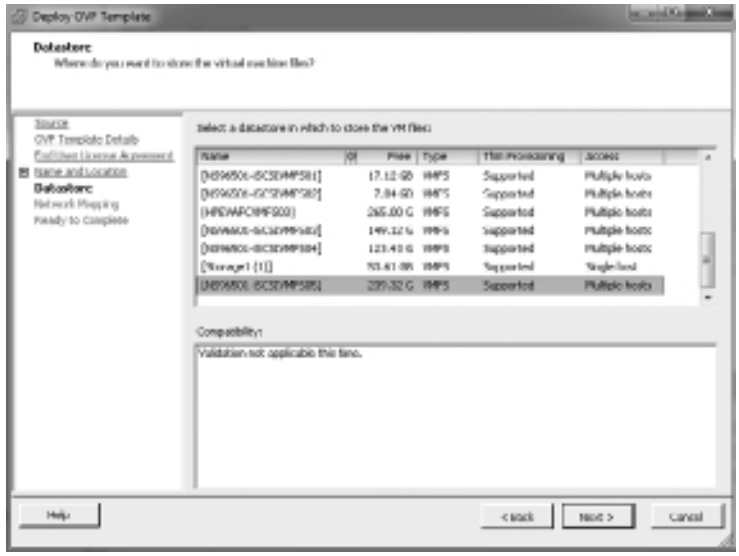

8. The Network Mapping screen, shown in Figure 5.74, is used to map the network used by the appliance to a usable network in your vSphere environment. This network will be used to communicate with AppSpeed probes. Select a network, and click Next.

FIGURE 5.74
Select a network to connect the virtual machine.

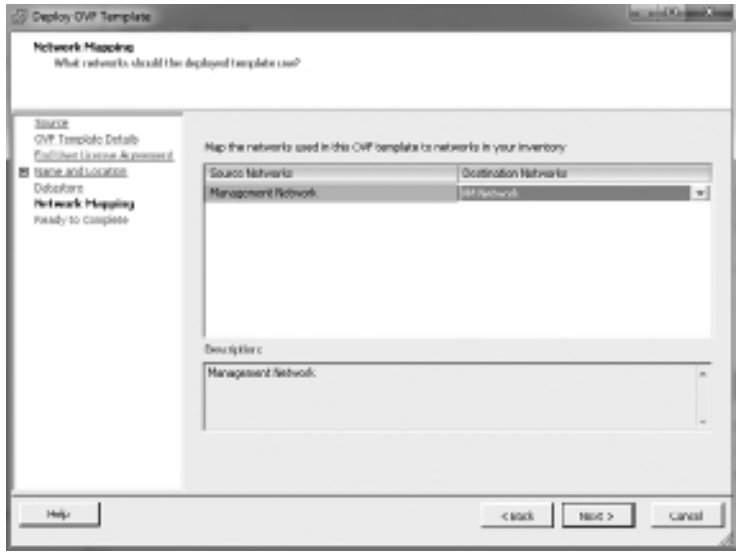

9. Review the Ready To Complete screen, shown in Figure 5.75, and verify that all the information is correct. During deployment notice the progress bar (Figure 5.76) and the deployment of the virtual disk (Figure 5.77). Once complete, the AppSpeed appliance will be ready for configuration.

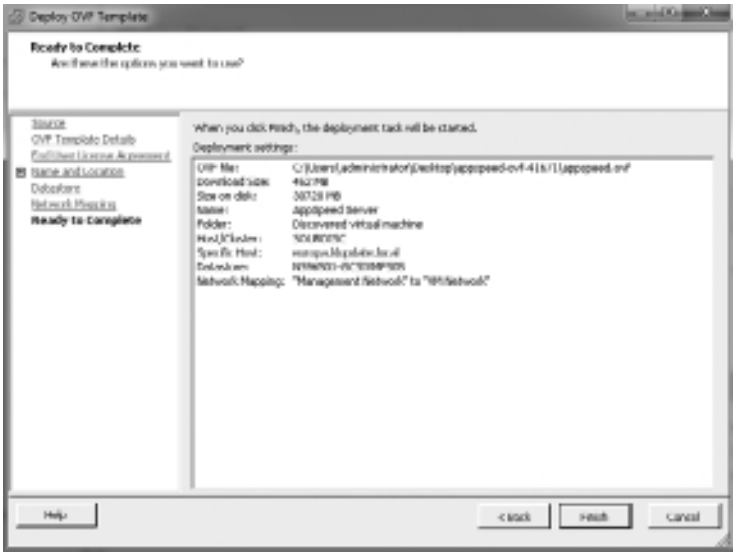

FIGURE 5.75
Review the summary, and click Next to deploy the AppSpeed server.

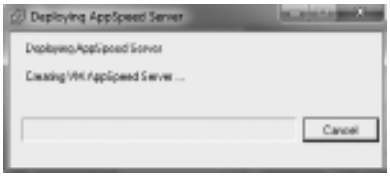

FIGURE 5.76
The status bar shows progress during deployment.

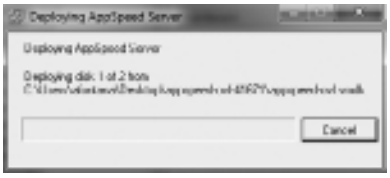

FIGURE 5.77
AppSpeed appliance virtual disk deployment

Once deployed, the AppSpeed Server appliance appears alongside other virtual machines. Figure 5.78 shows the default settings, version, and resource usage for the virtual machine on the Summary tab.

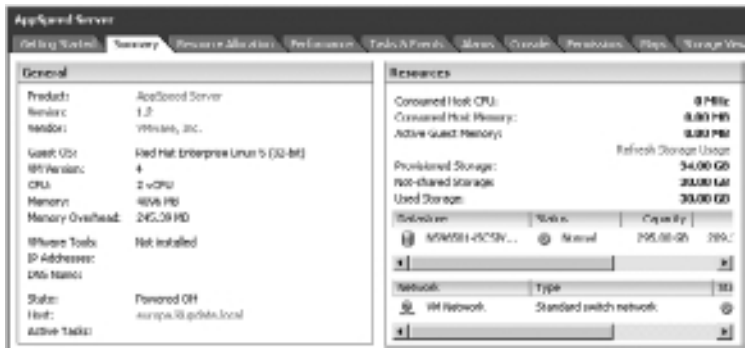

FIGURE 5.78
AppSpeed Server virtual machine summary

To configure the AppSpeed Server, follow these steps:

1. Power on the virtual machine.
2. Open a console or select the console tab, and enter **`setup`** at the login prompt.
3. Follow the on-screen instructions to set the network settings and time zone, associate the appliance with the vCenter server, and create the instance name.
4. Confirm the settings (Figure 5.79) by typing **yes**.

The appliance will reboot and will be ready for use.

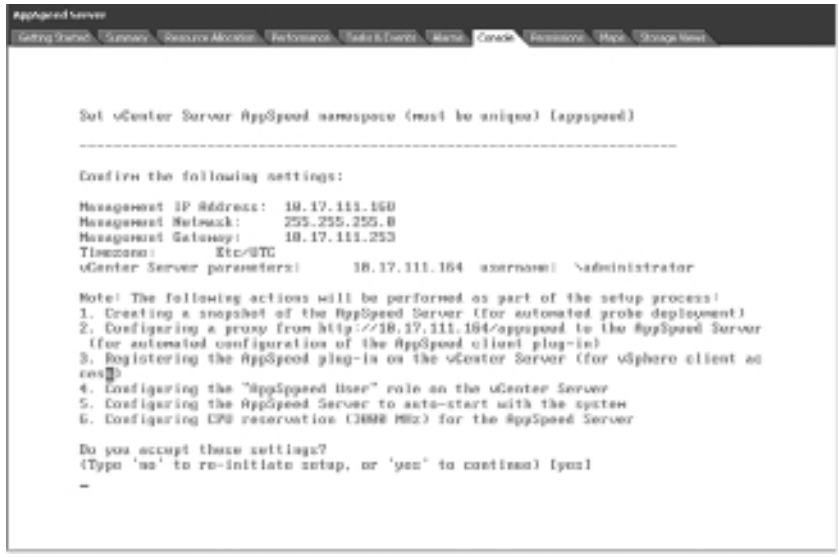
AppSpeed Server
Getting Started
Summary
Resource Allocation
Performance
Tasks & Events
Alarms
Console
Permissions
Maps
Storage Views
Set vCenter Server AppSpeed namespace (must be unique) [appspeed]
Confirm the following settings:
Management IP Address: 10.17.111.160
Management Network: 255.255.255.0
Management Gateway: 10.17.111.253
Timezone: Etc/UTC
vCenter Server parameters: 10.17.111.164 username: \administrator
Note! The following actions will be performed as part of the setup process!
1. Creating a snapshot of the AppSpeed Server (for automated probe deployment)
2. Configuring a proxy from http://10.17.111.164/appspeed to the AppSpeed Server
(for automated configuration of the AppSpeed client plug-in)
3. Registering the AppSpeed plug-in on the vCenter Server (for vSphere client ac
ess)
4. Configuring the "AppSpeed User" role on the vCenter Server
5. Configuring the AppSpeed Server to auto-start with the system
6. Configuring CPU reservation (3000 MHz) for the AppSpeed Server
Do you accept these settings?
(Type 'no' to re-initiate setup, or 'yes' to continue) [yes]

FIGURE 5.79
AppSpeed setting confirmation

The AppSpeed vCenter client plug-in provides the interface to AppSpeed. To begin using the features of AppSpeed, follow these steps:

1. Install the plug-in by selecting Plug-ins from the vCenter client and selecting Manage Plug-ins, as shown in Figure 5.80.
2. From the Plug-in Manager, locate the AppSpeed plug-in from the list of available plug-ins.
3. Click the link labeled Download And Install (Figure 5.81).

FIGURE 5.80
Launch the Plug-in Manager to view installed and available plug-ins.

FIGURE 5.81
Install the AppSpeed plug-in from the list of available plug-ins.

Once installed, the AppSpeed link will appear in the Solutions And Applications list on the vCenter client home screen, as shown in Figure 5.82.

FIGURE 5.82
AppSpeed link within the vCenter client home screen

Click the AppSpeed link to begin using the tool. If this is the first time the client is connecting to AppSpeed, you will have to provide the name of the AppSpeed instance (Figure 5.83). The instance name was created during the configuration of the AppSpeed server; the default is AppSpeed. Click Connect to continue.

FIGURE 5.83
Enter the name of the AppSpeed instance to connect to AppSpeed.

For AppSpeed to begin collecting data about your network of virtual machines and applications, you must deploy at least one AppSpeed probe. To deploy the probes, follow these steps:

1. Select the Administration tab within the AppSpeed application, and select Probes Management from the list on the left.

From this screen, you can deploy, delete, and power on or off your deployed probes. This screen also provides a quick look at the health and traffic throughput for specific probes.

2. Select a host, and click the Install Probe link, as shown in Figure 5.84.

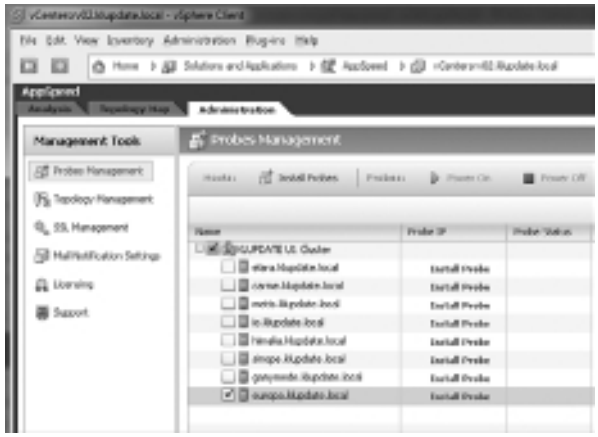

FIGURE 5.84
Deploy the AppSpeed probe to begin monitoring virtual machine network traffic.

The probe deployment screen, shown in Figure 5.85, prompts for the required information.

3. Provide a name for the probe, and specify the storage device to store the probe virtual machine files.

4. The management network will be used to communicate with the AppSpeed server. Provide an IP address or allow the probe to use DHCP to acquire an address.
5. The AppSpeed Probe can monitor traffic on all vSwitches available on the host or on a specific set of vSwitches. Select the vSwitches that you want to monitor traffic on, and click Install Probes.
6. The confirmation screen in Figure 5.86 will give a summary of the settings; click OK to begin the installation task.

FIGURE 5.85
AppSpeed probe installation screen

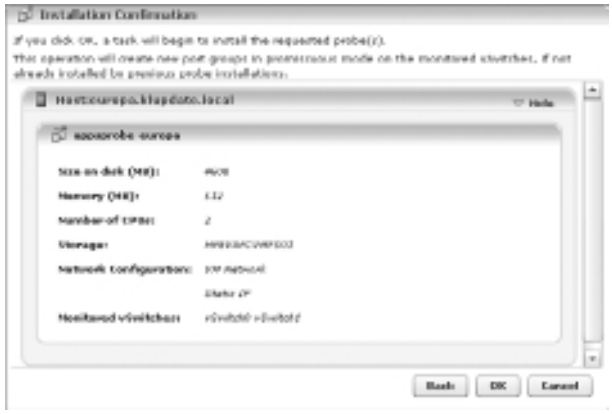

FIGURE 5.86
AppSpeed probe installation confirmation screen

Once installation has completed, the probe will appear on the Probes Management screen under the associated host. The status will appear as Running for a healthy, powered-on probe (Figure 5.87).

Traffic traversing the monitored vSwitches collected by the AppSpeed probes and processed by the AppSpeed server will begin populating the AppSpeed Analysis And Topology Map screen (Figure 5.88). AppSpeed will begin to provide valuable statistics such as latency, usage, and availability of your applications as well as topology maps and dependency information.

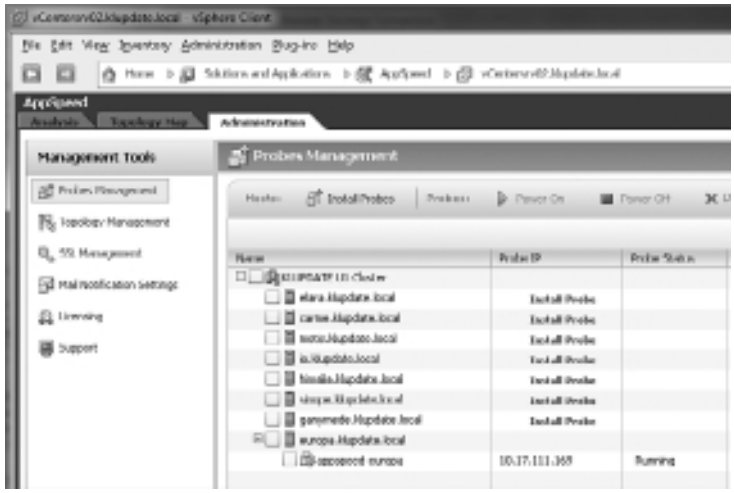

FIGURE 5.87
Healthy AppSpeed probe deployed and running

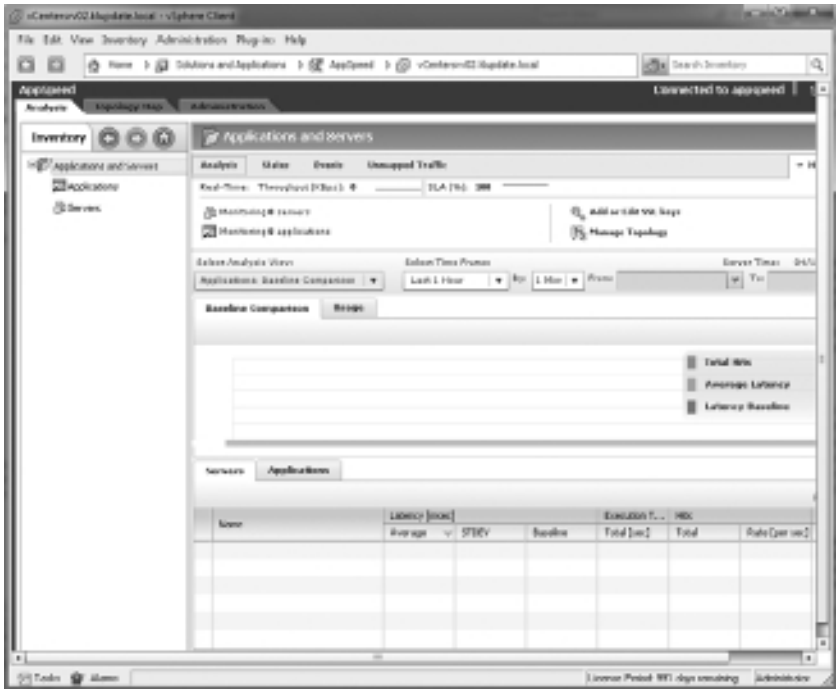

FIGURE 5.88
AppSpeed dashboard

Summary

IIS may not be the most dominant web server out there, but it's definitely gaining momentum. Windows 2008 introduced a completely rewritten version of Internet Information Server with version 7.0. New management tools have brought consistency to Microsoft's management consoles, improved delegation of authority, and improved command-line and script-based management. With the module-based approach for installing IIS and support for Windows Server Core, securing your IIS-based web servers has never been easier.

Virtualizing IIS is not difficult as long as you are aware of certain characteristics of your websites and applications. Taking the knowledge you've gained about the performance characteristics of IIS and having a reference for what counters to evaluate will help trend the usage and performance of your web servers and help you properly scale up or scale out your environment. When scaling out, you may want to take advantage of load balancing, and being aware of the configuration options you have when in a virtualized environment will ease the task of deploying Windows NLB. Deploying in a virtual environment not only gives you the ability to partition your hardware into more efficient units, but it also helps in securing the environment when you take advantage of optional components such as vShield zones. With IIS down, you can begin tackling some of the applications that will take advantage and complement IIS such as SharePoint and SQL Server.

Chapter 6

Virtualizing Exchange Server

Messaging and collaboration have become critical components in most if not all business operations. In businesses, email is integrated into most business processes, including the help desk, bug tracking, archiving, document management, and monitoring operations. When email is unavailable, business productivity decreases, and foot traffic in the break room increases dramatically. Escalations and calls to the help desk increase dramatically. There are many reasons for an Exchange server outage to include server hardware failure, database corruption, and performance degradation because of server storage over-subscription. Virtualizing Exchange Server can help mitigate most of the issues associated with a physical deployment by doing the following:

- Allowing you to right-size your Exchange Server deployment by adding resources as needed with just a reboot.
- Moving your Exchange Server to a more capable host to perform host upgrades without downtime to your end users through the use of VMotion.
- Providing simplified availability through the use of VMware HA if your Exchange Server virtual machines hang or you get a blue screen.
- Allowing you to create simplified disaster recovery scenarios without the requirement of similar hardware. A virtual machine can run on any VMware-supported physical server that has the ESX hypervisor installed.

Over the past 10 years, the Microsoft Exchange Server product team has managed to create a messaging and collaboration product that captures more than 50 percent of the messaging and collaboration market share all while dominating this space.

Microsoft has worked very hard on this release to make the product easier to design, deploy, manage, and recover. With all the new features and enhancements, Microsoft Exchange 2010 is threatening to dominate even more of the messaging and collaboration market. There are a lot of compelling reasons to transition from an Exchange Server 2003 or 2007 environment to an Exchange 2010 environment.

Exchange Server 2010 is totally virtualization friendly, more so than Exchange 2003 and Exchange 2007. We will briefly cover some of the more important features and reasons you might want to start planning for a virtualized Exchange 2007 or 2010 environment.

Exchange Server Virtualization Trends

Prior to 2007, not many medium to large organizations were even thinking about virtualizing Exchange Server, let alone actually doing it. These were the major objections to virtualizing this application:

- Lack of Microsoft support for the virtualized Exchange Server application
- Virtualization overhead
- Poor application performance when virtualized
- Poor storage I/O performance when virtualized
- The "If it ain't broke, don't fix it" mentality
- Fear of loss of employment

Some experimented with virtualizing Exchange Server 2003 but abandoned the effort when they experienced severe performance issues with Mailbox server virtualization.

Microsoft also didn't make it easy to virtualize its applications, which created some major barriers to the virtualization of tier-1 applications. In the fall of 2008, Microsoft announced the Server Virtualization Validation Program (SVVP) and, in doing so, took a major step forward into the virtualization journey.

With this program, Microsoft announced support for its major applications, and Exchange was one of the first. Before this announcement of support, some brave and daring customers took a modular approach and virtualized the bridgehead servers and front-end servers and left the back-end Mailbox servers on physical servers.

Despite the warnings from Microsoft that they would not be supported or that they would be required to reproduce the issue on physical hardware in order to receive support, customers continued to pursue virtualization.

With the maturity of the VMware virtualization platforms and the formation of the tier-1 applications performance team within VMware, VMware was able to extensively test Exchange Server 2003 and Exchange 2007 and produce recommendations and guidance for virtualizing the Exchange products. After producing these recommendations and guidance, VMware started to see an uptick in the virtualization of the Exchange Server product. Most of this virtualization was being done on Exchange Server 2007 with minimal efforts being applied to the Exchange Server 2003 platform.

VMware then started to work with the engineering folks to enlist the various server and storage manufacturing partners to improve the performance of virtualized tier-1 applications. Microsoft Exchange Server was one of the first tier-1 applications with which VMware did extensive performance testing and performance optimizations on the VMware virtualization platform.

To prove to the information technology world that the virtualization of Exchange was entirely possible, VMware took on the challenge of virtualizing its internal Exchange environment.

Two authors of this book, Alex Fontana and Charles Windom, were the principal architects for the design and implementation. During this process, the architects worked directly with VMware engineers, members of the file system group, and the high-availability folks within VMware to ensure that the virtualization project would meet or exceed the performance of the existing Exchange Server 2003 implementation.

TIP For more information on the VMware Internal Exchange virtualization design and implementation, please visit the VMware website at `www.vmware.com/solutions/business-critical-apps/exchange/resources.html`.

In an effort to improve the performance for the virtualized Exchange deployments, VMware started working with its server and storage partners to create real-world Exchange 2003 and 2007 configurations. Today, work is still being done to constantly improve the performance of the vSphere platform as it pertains to running the Microsoft 2010 suite of tier-1 applications.

Exchange Server Comparisons

As the Microsoft Exchange Server platform has continued to evolve over the years, it has changed from the typical all-in-one message server to a more distributed environment.

Exchange Server 2003

Exchange Server 2003 offered the following distributed server roles; some of these server roles were more difficult if not impossible to virtualize:

Front-end server This is responsible for external client connectivity (Outlook Web Access, POP3, IMAP, RPC/HTTP, SMTP, and Outlook Mobile Access, to name a few). This server role was easy to virtualize because it wasn't very demanding on storage resources, and the VMware platform provided more than adequate CPU, memory, and network resources for these roles to perform as well as a physical server.

Exchange bridgehead server This is responsible for delivering messages to Mailbox servers outside the client sender site. Exchange Server 2003 utilizes routing groups to send and receive messages between Exchange sites, such as a local Exchange site in Palo Alto and a remote site in Japan. This was another server role that could be virtualized, although you needed to pay more attention to the configuration of the virtual machine for this role because it required more CPU, memory, network, and storage resources than the Exchange front-end server role.

Back-end server This server role contains the database(s) used to hold the mail databases. This is also the server role that provides mailbox and public folder access. It can also (optionally) provide NNTP, POP, and IMAP4 access. This server role required careful design and planning when virtualizing.

Virtualizing this server role in large to very large organizations proved daunting if not impossible. Because of the demanding CPU and storage requirements for this role, these organizations would have to deploy more virtual machines and increase the management overhead associated with managing those virtual machines. This proved to be a major obstacle in virtualizing this server role.

Exchange Server 2007 and 2010 Server Roles

With the introduction of Exchange 2007, Microsoft changed the configuration and naming convention of the server roles. Instead of the typical three server roles, Microsoft introduced five new server roles; three of the five had somewhat identical functions to the previous server roles.

There are no longer the Exchange front-end, bridgehead, and back-end server roles. The new names of the roles going forward are as follows:

- Exchange Edge server role (EDGE)
- Exchange Client Access server role (CAS)
- Exchange Hub Transport server role (HUB)
- Exchange Mailbox server role (MBX)
- Exchange Unified Messaging server role (UM)

Each Exchange server role performs a specific function in the Exchange organization. For each role, only the Exchange server services get installed and started for that role. You no longer have to worry about disabling multiple Exchange services and deleting databases on server roles that do not require them.

In the next sections, we will give a brief explanation for each role's responsibility in the Exchange organization.

Edge Server Role

This server must be deployed at the perimeter of the network. It should not be a member of the domain.

Figure 6.1 highlights the Exchange 2007/2010 Edge server role. As you can see in the figure, there are two Edge servers deployed in the demilitarized zone (DMZ) of the network.

You can deploy multiple Edge servers in the perimeter network. When you do deploy multiple Edge servers, they will provide redundancy and failover capabilities for inbound email.

You also have the ability to load balance your SMTP traffic to multiple Edge servers by defining multiple mail exchanger (MX) records with the same priority in the DNS database for your mail domain.

The Edge server also provides messaging hygiene known as antispam and antivirus control for email incoming to your Exchange organization and outgoing to the Internet.

Edge transport servers can also prevent information that could be confidential to your organization from being sent outside the organization through the use of edge transport rules.

You can create and assign rules to control message flow out of the Exchange organization, and if a message meets the criteria of the created rule, that message will be blocked from being sent outside the organization.

The Edge transport server can also perform inbound and outbound SMTP address rewriting on messages. This can be especially useful to provide a consistent appearance of SMTP addresses in email organizations that perform lots of acquisitions.

Virtualizing this role is easy because it is not very demanding on CPU, memory, and disk resources. However, be mindful of your network resources when deploying multiple servers because this places additional demand on the physical network adapters in the ESX host. Ensure that you provide adequate network resources and redundancy through the use of ESX network teaming.

FIGURE 6.1
Edge server role

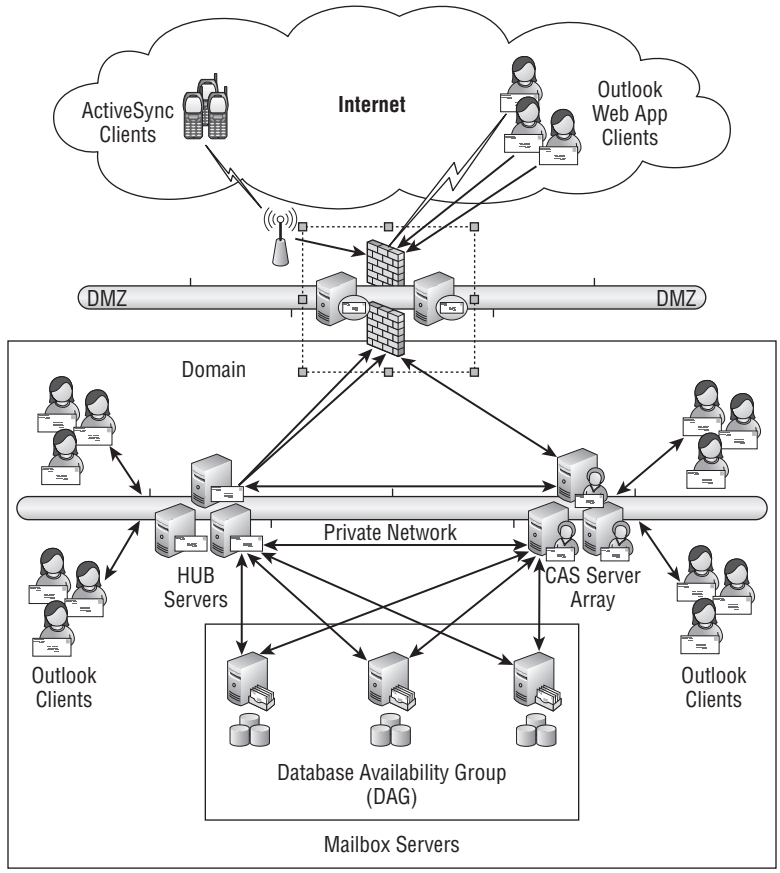

Client Access Server Role

The Client Access server (CAS) role hosts all client protocols such as Outlook Web Access (Outlook Web App for 2010), POP3, IMAP4, Outlook Mobile Access (OMA), and Outlook Anywhere (formerly known as RPC over HTTPS). The Client Access server role must be installed in every Active Directory site that has a Mailbox server installed.

Figure 6.2 shows the CAS role deployed as a CAS array on the internal network protected by the firewall.

The Client Access server is also responsible for providing access to the Exchange free/busy data and also allows Exchange 2007 and greater clients to download automatic configuration settings using the Autodiscover service.

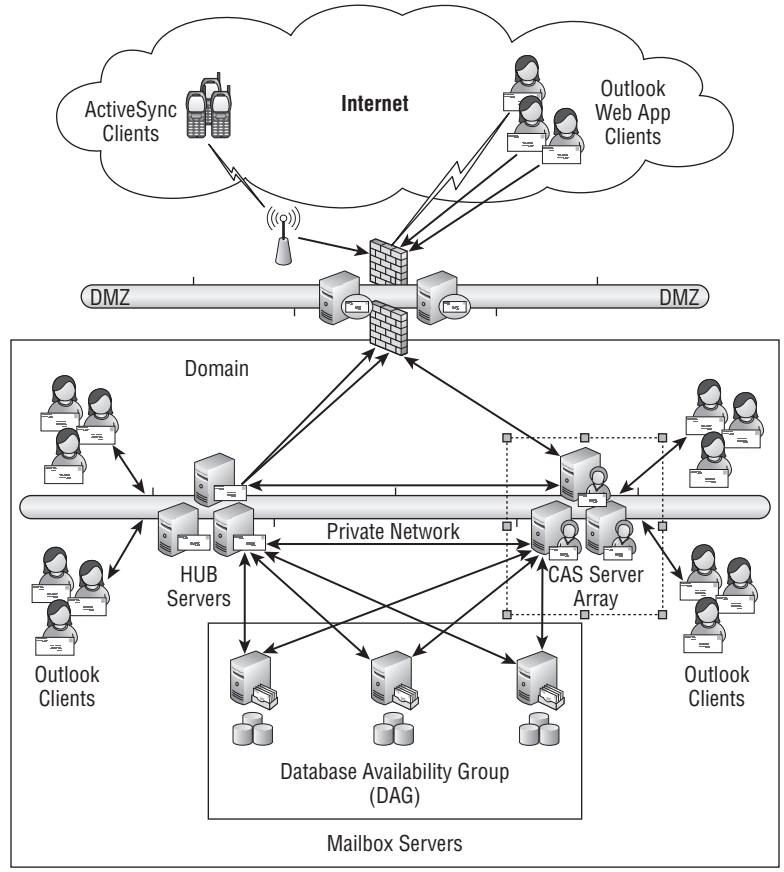

FIGURE 6.2
Client Access server array

To increase the scalability, availability, and redundancy of the Client Access server, in 2010 there were changes made to the Client Access server role. These are some of the fundamental changes:

- New client access server arrays (CAS arrays) provide availability and load balancing. The failure of a member of the CAS will automatically connect the client to another member of the CAS array.
- CAS arrays allow a higher number of concurrent connections per server and a higher number of mailboxes per server.
- The new RPC endpoint on the Client Access server role replaces the RPC endpoint on the Mailbox server role.
- New business logic routes Outlook clients to the CAS server to access a mailbox instead of logging onto the Mailbox server.

Exchange 2010 CAS server roles use RPC encryption by default, so Outlook 2007 and greater clients are compatible with Exchange 2010 CAS servers. Outlook 2003 clients will need to modify their settings to force RPC encryption in order to connect to a 2010 CAS server or array. Exchange Server or Windows Server administrators can also create a group policy to force Outlook 2003 to use RPC encryption.

This server role in Exchange 2007 can be easy to virtualize because you can also deploy multiple servers to provide load balancing and redundancy capabilities for the CAS server role. With Exchange Server 2010, be mindful that the CAS server role has taken on additional responsibilities (mentioned earlier in this section) and requires more server resources in the form of CPU, memory, disk, and network resources. The same rule of deploying multiple servers to provide load balancing and redundancy applies, but now you have the CAS array option. In actual production scenarios, we have seen that IMAP places additional demand on physical server resources than MAPI and could mean that you will have to deploy additional virtual machines to meet this additional demand.

TIP If your Outlook 2003 users cannot connect to an Exchange 2010 CAS server, or receive the error messages "Cannot start Microsoft Office Outlook. Unable to open the Office Window. The set of folders cannot be opened or unable to open your default email folders. The information store cannot be opened," then try enabling the Encrypt Data Between Microsoft Office Outlook And Microsoft Exchange Server option in the Outlook 2003 security settings.

Hub Transport Role

The Hub Transport role is responsible for moving messages through the Exchange organization. When we say "through the Exchange organization," we mean inside and outside the Exchange organization.

The Hub Transport role uses Active Directory sites to route email messages to local and remote Exchange sites in the email organization. Deploying multiple Hub transport servers in the same Active Directory site will provide availability as well as load balancing. There must be at least one Hub Transport server installed in every Active Directory site that has the Mailbox server role installed.

In the absence of the Edge server role, the Hub Transport can be configured to send and receive messages to and from the Internet as well as provide messaging hygiene, messaging policy, and compliance using compliance and transport rules.

Enhancements to the Hub Transport server role in Exchange Server 2010 include the following:

- Shadow redundancy provides redundancy for messages the entire time that they are in transit.
- Write I/O is reduced by about 50 percent.
- The transport server becomes stateless.
- The Hub Transport dumpster has been enhanced to use a database replication feedback mechanism to control what messages remain in the dumpster.
- Messages are deleted from the dumpster once it has been replicated to all database copies.
- Larger messages are supported without triggering back pressure.

In Exchange Server 2007, this server role is fairly simple to virtualize, but you must take into consideration the disk I/O requirements for the Hub Transport server, because it maintains the queue database. We recommend using RAID 1 or RAID in larger environments. With Exchange Server 2010, this server role becomes easier to virtualize because of the reduction in disk I/O and the stateless database. It is safe to utilize RAID 5 for the database, and we recommend that you deploy at least two Hub Transport servers in every site that hosts an Exchange server.

Figure 6.3 shows a deployment of three Hub Transport servers on the internal or private network. This will facilitate load balancing as well as redundancy for the Hub Transport server.

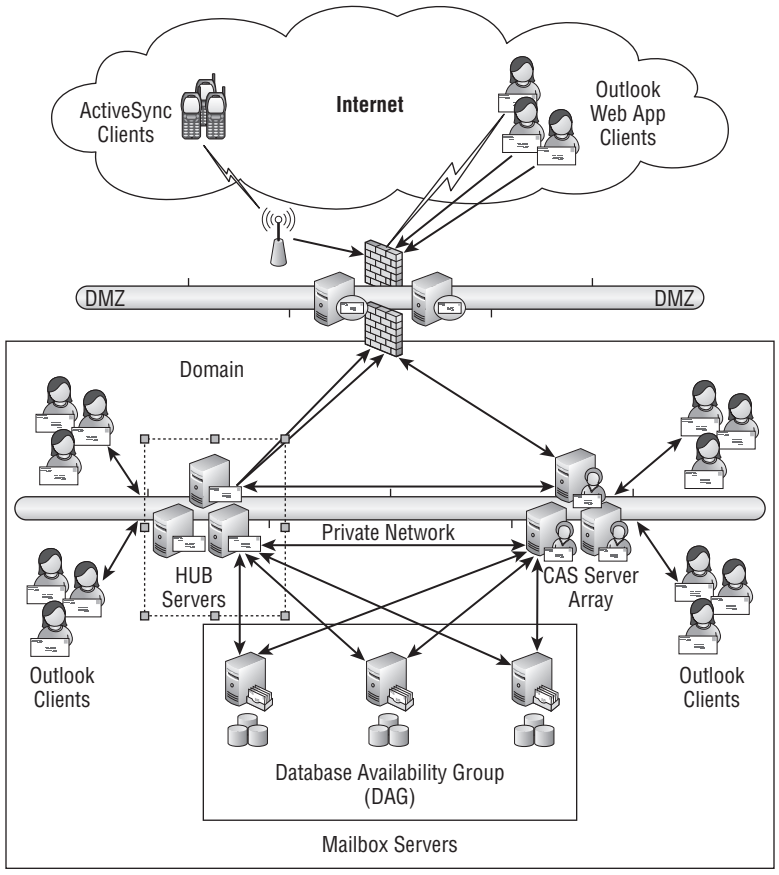

FIGURE 6.3 Exchange 2007/2010 Hub Transport server role

Mailbox Server Role

This server role will be the heavy lifter of your Exchange 2010 deployment. Its sole responsibility is to play host to the mailbox and public folder databases. In Exchange 2007 and 2010 Server editions, the public folder is optional and has been provided for backward compatibility.

All connections with the exception of the public folders are accomplished through the Client Access server role. Connections to the public folder databases are still handled by the RPC

endpoint on the Mailbox server role, so clients must still connect directly to the Mailbox server role to access the public folders.

Figure 6.4 shows three Exchange 2010 Mailbox servers deployed as a database availability group (DAG). We will get into DAGs later in the chapter.

FIGURE 6.4
Exchange Server 2010 Mailbox servers

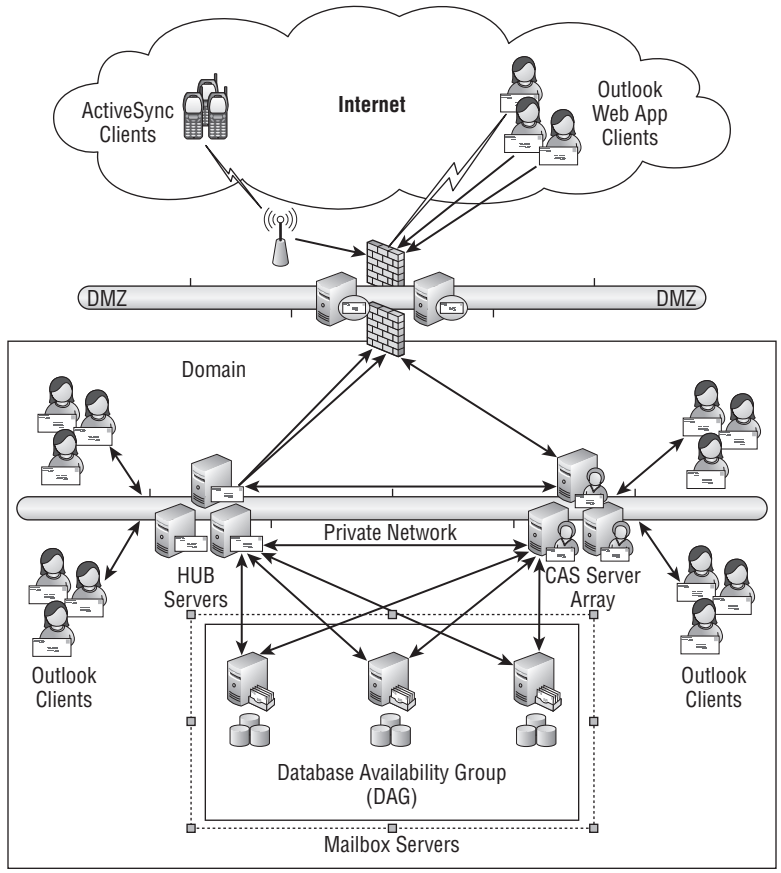

The Mailbox server role is also the logical unit in a high-availability configuration. Because the Mailbox role performs the database hosting role only, you must have installed a Client Access server and a Hub Transport in any Active Directory site that will host Exchange messaging services.

As mentioned earlier, the CAS role will handle all client connections for mailbox access, and the Hub Transport role will route message throughout the Exchange organization.

New for the Exchange 2010 platform, the concept of the storage group has been deprecated. The term database has been redefined to mean an Exchange database and its transaction log files; they are now thought of as a single unit. In the Exchange 2010 Enterprise edition, each Mailbox server role can host up to 100 databases.

The Exchange 2010 Standard edition can host only five databases per server, but now you can also use DAGs on the Standard edition when installed on Windows Server 2008 Enterprise. It is important to note that Exchange Server 2010 Standard edition can upgraded to Enterprise edition without reinstalling the product. Because DAGs use Windows clustering services, Windows Server 2008 Enterprise edition is required for implementing DAGs.

Exchange databases are no longer pinned to the Exchange server. Because of this fundamental configuration change, all databases created on the Exchange Mailbox role must have a unique name.

Real World Scenario

Real-World Lesson Learned

When first testing the Exchange 2010 platform with the new LoadGen 2010 tool, a colleague of ours created four individual mailbox databases and then copied the databases and renamed them manually to create a total of twelve databases for testing. When he tried to run the LoadGen 2010 tool, he experienced many errors, one of which was a feedback loop.

We were asked to troubleshoot the configuration and after googling for the specific error, we decided to dismount the last eight databases of the twelve and rerun the LoadGen test tool.

To our surprise, all of the errors disappeared, and we were able to run a successful test. We then deleted the last eight databases, created eight new mailbox databases, and proceeded to run the tool again. All tests completed successfully.

This was a lesson learned for our colleague as well as ourselves. Exchange databases must be truly unique, or you will experience undocumented issues. There was no information anywhere regarding this error message.

Unified Messaging Server Role

This provides the ability to access mail messages via phone or computer. This role is not currently supported under virtualization. Although there are some customers actively virtualizing the Unified Messaging (UM) server role with success, Microsoft explicitly states that it will not support this role virtualized and recommends that you deploy this role on physical servers. You should install the UM server role in an Active Directory site after the Hub Transport role is installed.

Not much testing has been done internally to prove whether this role is a viable candidate for virtualization. Apart from this, we will discuss only some of the capabilities of this role under the Exchange 2010 platform. These are some of the features and enhancements for Exchange 2010:

Auto Attendant Allows for the creation of custom menu prompts to navigate the UM system. The user can then use their voice or touchtone keypad to navigate the system.

Call Waiting light Alerts user of waiting messages.

Subscriber Access Allows a user access to their mailbox, calendar items, and messages using a regular phone line.

Text-to-Speech Allows user to listen to messages in their Outlook mailbox.

To date Microsoft does not support the Unified Messaging server role virtualized because of heavy demands on server CPU and memory. With the introduction of vSphere, however, you can have an eight-vCPU 255 GB memory virtual machine. With these new virtual machine capabilities, some customers have started to virtualize this role. We recommend that you follow the Microsoft Support guidelines about virtualizing Exchange Server 2010 to ensure that you get the best support that Microsoft has offer for virtualized Exchange deployments.

These roles could be deployed on a single server or on multiple servers for modularity. Even though you could install your Exchange 2003 server roles onto multiple servers, you would still get all the Exchange Server services installed and started on each of the server roles. This became a real point of frustration with Exchange and Windows server administrators when trying to secure the various server roles.

Exchange Server 2003 and prior versions were all 32-bit environments, and the maximum addressable memory that Exchange Server 2003 and prior versions could handle was 4 GB. You could increase this limit by using various `boot.ini` switches, but under heavy usage scenarios, you could exhaust the Windows kernel resources and fragment the usable memory address space. The only way to recover from these issues was to reboot the server. Table 6.1 describes some of the differences among the various Exchange Server versions since Exchange Server 2003.

TABLE 6.1: Comparison of Exchange Server Versions

EXCHANGE SERVER 2003	EXCHANGE SERVER 2007	EXCHANGE SERVER 2010
32-bit environment.	64-bit environment. 32-bit version (not supported in production).	64-bit environment.
4 GB max addressable memory.	Multiterabyte max addressable memory.	Multiterabyte max addressable memory.
About 932 GB addressable database cache.	Multigigabyte database cache for better performance.	Multigigabyte database cache for better performance.
4 storage groups	50 storage groups	Storage groups removed from Exchange 2010. Databases and logs still remain, but both are now referred to as just a database. 100 databases per server max.

TABLE 6.1: Comparison of Exchange Server Versions *(CONTINUED)*

EXCHANGE SERVER 2003	EXCHANGE SERVER 2007	EXCHANGE SERVER 2010
Storage I/O very heavy because of small block size (about 4 KB) and database cache. Very difficult to scale beyond 2,500 users per Mailbox server and maintain large mailboxes and good end-user experience.	Larger database cache, increased block size and write coalescing to increase performance, and reduction in storage I/O to about 50 to 70 percent.	Schema changes along with sequential database access and large database cache allow for even more storage I/O reduction (around an additional 50 to 70 percent) and larger mailboxes.
Administrative groups used to delegate control and access to Exchange organization. Routing groups used to send and receive mail to nonlocal Active Directory sites.	Routing changes in Exchange 2007 and the introduction of the Hub Transport server role. With Exchange Server 2007, routing is now performed using Active Directory sites.	Routing remains pretty much the same with enhancements to the Hub Transport server role. Shadow messaging is introduced. Shadow messaging ensures that messages are delivered to each hop in the delivery path. If delivery fails anywhere in the delivery path, the message is re-sent from the previous hop to where the message failed. Role-Based Access Control (RBAC) is introduced in Exchange 2010. RBAC provides the ability to granularly grant access permissions to the Exchange 2010 organization.
	Active Directory security groups are used to grant access and delegate access to the Exchange 2007 organization.	The database page size has been increased to 32 KB.
		Database mobility means that databases are no longer pinned to a specific server.
		DAGs allow you to maintain multiple copies of your Exchange databases and provide immediate failover in the case of server or database corruption.
		Microsoft removed the server as the dependency for the Exchange Server databases. Databases have been moved to the organization level of the Exchange Server.

These are just some of differences that make Exchange 2007 and 2010 extremely virtualization friendly. There are whole hosts of other enhancement and improvements, and we invite you to examine these features and improvements to the Exchange 2007 and 2010 platforms.

TIP You can find out more about Exchange Server on the Microsoft Exchange Server home page at `www.microsoft.com/exchange/`.

Exchange High Availability

Exchange Server high-availability options and features have changed with the evolution of the Exchange product. Microsoft went from stand-alone servers to two-node clustering, and over time and new releases of the Exchange product, it offered four-node clustering.

Exchange Server 2003 and the Windows Server product increased the clustered node numbers to eight nodes, with a seven-active-and-one-passive clustered configuration. Although clustering could provide server and service availability, it was still not considered by the industry as a true highly available solution.

In a true highly available solution, you will find the following characteristics:

- Server or hardware availability
- Service availability
- Data redundancy and availability

Hardware Availability

In clustered solutions, you will encounter a multiple of servers, usually configured identically, all working together to increase the uptime and availability of the application or infrastructure services configured to run as a cluster service or application.

In these scenarios, there is usually a resource DLL that is registered as a cluster resource upon installation of the infrastructure service or application unless the application is a generic application that you are trying to cluster. In this case, you are creating a generic clustered service or application.

In the event of hardware failure or outage, the remaining nodes of the cluster will provide service availability by performing a failover to one of the surviving nodes of the cluster. In most large Exchange Server failover clustering scenarios, it is typical to find 6+2, 5+3, and 4+4 active-passive scenarios.

Service Availability

In the area of service availability, you will want to ensure that in the face of some business disruptive event, the Exchange Server services will be able to provide continued services to the end users without extended downtimes.

Through the use of the clustered resource, the service or services will be constantly monitored for availability, and in the event of some business disruptive event, the service or services will be started on a surviving node of the failover cluster.

In some clustered application scenarios, an additional set of identical services is already running on additional nodes within the cluster. In this scenario, the actual application or infrastructure service recovery time is shorter.

You encounter this scenario in the Exchange Server 2007 clustered continuous replication (CCR) configuration. A subset of the services is already running on the second clustered node, providing the replication services as well as playing replicated transaction logs in the second copy of the Exchange Server database or databases.

Data Redundancy and Availability

In any true highly available infrastructure or application solution, there must be data redundancy. Your clustered solutions are useless if you or your end users cannot access their data (their data being their mailboxes). Data redundancy is one area where traditional failover clustering actually fails.

In the traditional failover clustering scenario, you have multiple clustered nodes accessing a single copy of data stored on shared storage. This is known as a single copy cluster (SCC).

In this scenario, a single node of the cluster is in control of the shared storage at any time. When a node of the cluster fails or becomes unavailable, control of the data on the shared storage is transferred to a surviving node in the clustered scenario.

In the event of logical corruption, the data itself will become unavailable to end users, and recovery procedures will need to be initiated. This could be a simple volume shadow restore that in most cases could take anywhere from minutes to hours depending on the nature of the logical corruption.

This also brings to light the fact that if your clustered solution is locally deployed or deployed using a data center service model, where all nodes of the cluster are local to a single data center, failure of the data center will cause a complete cluster outage rendering services unavailable to all users.

With the introduction of Exchange Server 2007, Microsoft introduced a clustered solution known as clustered continuous replication (affectionately known as CCR). In the CCR scenario, you have a two-node majority node set (MNS) cluster. We will not go into great detail on MNS because it is outside the scope of this book.

With CCR, you have an active and passive cluster node that does not utilize a shared clustered disk model. Quorum is achieved through local disks on each clustered node as well as a file share witness (FSW). The file share witness is a file share usually placed on the Exchange Server 2007/2010 Hub Transport, although this file share could actually be placed on any server.

TIP Microsoft recommends that the FSW be placed on the Exchange Server Hub Transport server and have redundant network connectivity to all clustered nodes.

In the CCR scenario, you have two copies of the Exchange Server databases that are kept in synchronization through asynchronous replication provided at the Exchange Server 2007 application level. Through the use of CCR, you have two copies of the Exchange mailbox data.

This fulfills the data redundancy requirement for a true highly available Exchange 2007 deployment. But wait a minute, what about the data availability requirement for a true highly available Exchange Server deployment?

Exchange Server 2007 CCR allows for deployment in long-distance scenarios known as a *site resilience model*. In this model, we deploy one CCR node in one data center and the other node at a remote data center some distance way. Some of the shortcomings of this solution are network bandwidth and network latencies.

You must ensure that your network can support the replication traffic for the implementation of long-distance CCR. Although Exchange transaction logs have been reduced to 1 MB in size, you could have tens of thousands of them traversing your network over the average workday.

Microsoft also provided standby continuous replication (SCR) as a means for fast disaster recovery. In SCR, you replicated transaction logs to a remote site and delayed the playback of replicated transaction logs into the Exchange Server databases.

Enter database availability groups. In Exchange Server 2010, Microsoft introduced a more mature implementation of CCR and SCR. We like to call it SuperCCR, or CCR on steroids.

This creates a total mobility scenario for the Exchange databases. Exchange Server 2010 databases are now Active Directory objects. In the DAG scenario, copies of the Exchange Server 2010 databases are created on servers within a majority node set cluster now known as a DAG for Exchange Server 2010.

Windows Failover Clustering (WFC) functions have been totally hidden from the Exchange administrator. The Exchange administrator creates the DAG and adds Exchange Server 2010 servers to the DAG as members. In all actuality, you are creating a Windows Server majority node set failover cluster shell, if you will, and then adding nodes to the cluster.

There is now a new Exchange Server 2010 service that manages state of the active and passive copies of the databases. This is the Exchange Server 2010 Active Manager service, which runs on all Exchange Server 2010 servers that are participating in a DAG.

There are two incarnations of the Active Manager, the primary Active Manager or the (PAM) and the standby Active Manager (SAM).

The PAM runs on the Exchange Server 2010 hosting the active copies of the Exchange databases. It constantly monitors Exchange databases and can detect failures. When a failure is detected, the Exchange Server hosting the active databases asks the PAM to perform a failover if the server is still online.

If the server hosting the PAM is unavailable, the PAM role will be transferred to another Exchange server, and the failed databases will be mounted. The PAM also owns the cluster quorum resource and is responsible for making decisions about topologies changes within the DAG. The PAM updates the RPC Client Access service with details about the newly mounted databases so that client connections can be redirected to the server that now hosts the active database.

The SAM role assumes the standby role of monitoring the state of the Exchange database copies and receives requests from the PAM to activate healthy copies of the Exchange databases in the event of a mounted database or server failure.

TIP For more information on the Exchange Server 2010 Active Manager, please see the Exchange Server 2010 online documentation on the TechNet website (`http://technet.microsoft.com/en-us/library/dd776123.aspx`).

Workload Considerations When Virtualizing Exchange Server Roles

Virtualization has made it increasingly simple to deploy systems quickly and with very little effort. Unfortunately, this can also lead to very little thought being put into the performance characteristics of the systems you are deploying. Your virtual data center has a pool of processing power, memory, disk, and network. It's very easy to carve that up to suit the needs of your virtual machine, but what you may not be aware of is the number of resources you are tapping into that are shared by other workloads. This is why proper planning, testing, and tuning are especially important in a virtual environment.

In the physical world, you are able to isolate workloads, at least to a single physical server, if not across the network and storage layers. Although this ensures that you are getting the dedicated resources to the applications being hosted on that single piece of hardware, it leads to a large consumption of power, cooling requirements, and under-utilized hardware. This is largely why the industry is moving toward deploying more and more apps to virtualized hardware. How do we do so in a way that we can guarantee our applications will be as responsive as they were on their over-provisioned hardware?

- Know the performance characteristics of your application.
- Gather performance data from the current production environment; if data is not available, use recommended "building blocks" as a starting point.
- Inventory your current resource consumption to ensure there is proper capacity to move forward and to ensure you are not excessively over-committing your resources.
- Test using expected workloads in stage or preproduction environment.
- Provide adequate time to tweak virtual hardware if needed.

With this data in hand, you will be much better prepared to begin planning the production environment, and you will have a sense of how your environment will perform. As Exchange has evolved over the past few years, functions have been moved around, and requirements have changed dramatically. In the early days of Exchange, we required bridgehead servers to send and receive mail. That requirement was somewhat lifted with Exchange 2000 and 2003 as the SMTP service was incorporated into the Mailbox servers. Now with Exchange 2007 and 2010, we again have the concept of a bridgehead server in the form of the Edge and Hub Transport roles. This is just one of many changes that have occurred in the past few years. Needless to say, if you haven't touched Exchange since 2000 or 2003 (or even 5.5), chances are you are unfamiliar with the new performance characteristics of the latest versions. This section will help introduce you to some of the basic fundamentals for appropriately sizing each role individually and for combining some of these roles for smaller environments.

Edge Transport

As you learned in the previous section, the Edge Transport role is designed to be the entry and exit points for all mails entering your organization. This role is deployed on the perimeter network and provides a secure mechanism for providing SMTP relay and smart host services by presenting a minimal attack surface and not participating in your organization's Active

Directory forest. The Edge Transport role is the ideal location for scanning, filtering, and tagging messages based on policies set by your organization. This is done by employing transport agents that act on messages based on content, attachments, or virus and spam signatures.

Processor

The Edge Transport role is able to take advantage of a multiprocessor configuration of up to 12 processor cores. This is the maximum recommended configuration for the Edge Transport role. The maximum recommended configuration is not a maximum supported configuration; rather, it's a balance between performance and cost. In most environments, four to eight processors will be the recommended configuration to be used by the Edge Transport because most people will choose to scale out their environments to multiple smaller virtual machines rather than scale up. Scaling out allows for greater redundancy and can lead to greater performance in the virtual environment because of being able to spread the load across more vSphere hosts.

When planning processor configurations for virtualized Edge Transport servers, there are some key factors that will contribute to processor utilization:

Message rate How many messages must you be able to process per second?

Message size What is the average size of messages that are processed?

Transport agents What types of transport agents will be installed and enabled? This includes antivirus, spam, content filters, and so on.

The Edge Transport role is more difficult to size than the Hub Transport role because of the level of message hygiene and inspection that happens at this layer. Being the point of entry into most organizations, chances are this role will handle scanning for spam and viruses and will validate message recipients. This means that every message needs to be evaluated and acted upon. Table 6.2 provides a baseline that you can use to compare against your organization.

Table 6.2: Edge Transport Performance Metrics: Internal Microsoft Deployment

Performance Metric	Value
SMTP Connections/Sec	55
% Connections Accepted	80 percent
SMTP Messages IMF Scanned/Sec	3.7
% SMTP Messages Passed IMF Scanning	80 percent
SMTP Messages A/V Scanned/Sec	3
Avg. Message Size	70 KB
CPU Utilization	20 percent

Microsoft TechNet: `http://technet.microsoft.com/en-us/library/dd346701.aspx#EdgeMemory`

Memory utilization for the Edge Transport role is minimal, relatively speaking. Significant memory use by the Edge Transport role is usually attributed to the database cache for the queue database, scenarios involving large queues, and EdgeSync requirements. The recommended configuration for memory allocation to the Exchange 2007 Edge Transport role is 1 GB of memory per every virtual CPU, with a minimum requirement of 2 GB. Following this recommendation, an Edge Transport virtual machine with one virtual CPU would require at least 2 GB of memory and every additional CPU allocated above has a recommendation of one additional gigabyte or memory per CPU. Exchange 2010 has increased the minimum requirement to 4 GB and maintains the 1 GB per CPU recommendation.

A factor to consider when dealing with large organizations or when planning for large queue scenarios is the amount of memory that will be required to hold each message in memory. Since queue messages and recipient information are held in memory, you need to be sure you are adequately sized to deal with the large queue. Table 6.3 provides the memory overhead used to calculate potential memory usage for Exchange 2007 and can be used as a baseline for Exchange 2010 transport servers.

TABLE 6.3: Memory Overhead Factor

MEMORY FACTOR	MEMORY ALLOCATED
Per message	3KB
Per recipient	1KB

Microsoft TechNet: `http://technet.microsoft.com/en-us/library/bb738124(EXCHG.80).aspx`

STORAGE

Although not as storage intensive as the Mailbox role, the Edge Transport role needs to be able to provide a certain level of I/O responsiveness to perform well. When deploying on a virtualized platform, it is easy to take for granted the importance of the underlying storage design. To the vSphere administrator, a VMFS partition with a few hundred gigabytes of free space may seem like the ideal location for your Edge Transport's storage, but let the mail start flowing, and you may realize that the underlying physical storage design is not up to the task. Designing the correct storage environment for your virtual Edge Transport servers is no different from how you would design for a physical environment. You need to understand the I/O and space requirements to sustain the unexpected queue buildup that can occur from time to time.

The amount of space required mostly depends on how large we expect the queue database to grow and how much logging we want to enable and keep. The following should be considered when planning how much space to allocate to an Edge Transport server; this does not include space for the operating system and page file:

Message tracking logs The size will depend on the amount of mail processed by your organization. If you already have an Exchange deployment, you can determine the current log generation rate. How much data is kept is configurable by setting the number of day's worth of logs to keep.

Protocol logs Sizes vary based on mail activity but can be configured with a maximum directory size and maximum age.

Connectivity logs Same as protocol logs.

Transport agent logs Same as protocol logs but may vary depending on the agent.

Queue database The amount of space taken up by the queue database is determined by the average message size multiplied by the maximum queue allowable. As an example, if you are designing the Edge Transport to hold a maximum queue of 100,000 messages at 50 KB per message, you must allocate 5 GB to the queue database.

Transaction logs Space is generally not a concern with the Edge Transport transaction logs because the depth is limited by the use of circular logging.

The Edge Transport role must have at least 500 MB (Exchange 2007 SP1) of free space (4 GB if you are planning on running Exchange Server 2007 RTM) to avoid having Exchange activate the back pressure feature.

Exchange 2010 has added some intelligence into the free disk space requirements for the queue database and transaction logs for transport-related ESE databases. In Exchange 2010, back pressure monitors a percentage based on the disk size since a static number like 5 percent is relative to the disk space available. The formulas are as follows:

- Free hard disk space for message queue database

  100 * (hard disk size – fixed constant) ÷ hard disk size

  Fixed constant = 500 MB

  The result is represented as a percentage of consumed hard disk space and is the high level. These results are rounded down to the nearest integer. The medium level, by default, is the high level minus 2 percent, and the normal level is the high level minus 4 percent.

- Free hard disk space for message queue database transaction logs

  100 * (hard disk size – MAX(5 GB or 3 * DatabaseCheckPointDepth)) ÷ hard disk size

  The DatabaseCheckPointDepth parameter is defined in the `EdgeTransport.exe.config` file located in the `<Install Dir>\Bin` directory.

  The result is represented as a percentage of consumed hard disk space and is the high level. These results are rounded down to the nearest integer. The medium level, by default, is the high level minus 2 percent, and the normal level is the high level minus 4 percent.

Back Pressure

Back pressure is a monitoring feature introduced in Exchange 2007 that allows the system to monitor its resources, including disk space, memory, and CPU, and make temporary changes to allow for the excessive resource consumption to pass, usually a large queue. Exchange 2007 reacts to back pressure by rejecting incoming connections. Exchange 2010 accepts all incoming connections, but messages over those connections are accepted at a slow rate or rejected until resource pressure is relieved.

Performance of the storage provided to the Edge Transport role is essential to ensure efficient mail flow in and out of your organization. Understanding the I/O characteristics of the Edge Transport role is key to designing a proper storage solution.

The Edge Transport role will attempt to store incoming mail in memory and in the transaction logs to minimize writing to disk. However, if memory is low, only a portion of each message will be stored in memory and in the transaction logs; the rest of the message must be written to the database. For this reason, it is important to keep the database on properly sized and performing disks, whether you choose virtual disks or raw device disks (pass-through).

Along with the queue database, the temp directory must have well-performing disks and is usually placed on the same volume as the queue database. The temp directory is used during the content conversion process, as messages come in that require conversion, the data is streamed to the temp directory to be processed. The location of the queue database and temp directory is stored in the `EdgeTransport.exe.config` file in the Exchange install directory.

To change the default location of the queue database and temp directory, follow these steps:

1. Stop the transport service by issuing the following command: `net stop MSExchangeTransport`.
2. Navigate to the `%programfiles%\Microsoft\Exchange Server\Bin` directory (`%programfiles%\Microsoft\Exchange Server\v14\Bin in Exchange 2010`).
3. Locate and open the `EdgeTransport.exe.config` file in a text editor such as Notepad.
4. The `<appsettings>` section in the `EdgeTransport.exe.config` file contains the configurable settings of the edge transport process. The tags are as follows:

   `QueueDatabasePath`: Queue database location

   `QueueDatabaseLoggingPath`: Queue database transaction logs

   `TemporaryStoragePath`: Temp directory

5. Start the transport service by issuing this command: `net start MSExchangeTransport`.

```
<configuration>
<runtime>
<gcServer enabled="true" />
</runtime>
<appSettings>
<add key="QueueDatabasePath" value = "C:\Program Files\Microsoft\Exchange
Server\Queue" />
<add key="TemporaryStoragePath" value = "C:\Temp" />
<add key="QueueDatabaseLoggingPath" value = "C:\Program Files\Microsoft\
Exchange Server\Queue" />
...
</appSettings>
</configuration>
```

When designing the storage system for your Edge Transport servers, you will need a baseline I/O profile to work with. If you are unable to use data from your production environment,

you can start with numbers published on Microsoft TechNet's website (`http://technet.microsoft.com/en-us/library/bb738141(EXCHG.80).aspx`). Although originally published for Exchange 2007, they provide a good baseline when the expected load is unknown. These numbers include an average message size and the message rate. For the example in Table 6.4, a message size of 60KB is used with an expected message rate of 20 messages per second.

TABLE 6.4: Disk IOPS per Message (60 KB)

EDGE TRANSPORT DATABASE I/O	APPROXIMATE EDGE TRANSPORT I/O
Log write I/Os per message (sequential)	11
DB write I/Os per message (random)	4.5
DB read I/Os per message (random)	2.5
Total I/Os per message	18

Microsoft TechNet: `http://technet.microsoft.com/en-us/library/bb738141(EXCHG.80).aspx`

Using Table 6.4 and the message rate (20 messages per second), you can calculate the database and transactional I/O required to provide the performance necessary. In the following formulas, 20 percent has been added to provide for growth and traffic spikes.

DB IOPS per message = $(20 \times (4.5 + 2.5)) + 20\% = 168$

LOG IOPS per message = $(20 \times 11) + 20\% = 264$

As you can see, the required database IOPS came in at 168, with the log IOPS at 264. With the I/O and sizing requirements in hand, you can work with your storage team to design the best disk layout for your Edge Transport virtual machines.

Client Access Server

The Client Access server role has evolved from what used to be called the front-end server in Exchange 2003. At that time, the front-end server was nothing more than a landing spot for client connections that were then proxied back to the Mailbox server. This meant very little resources were required on these front-end servers since all the heavy lifting was still being done by the Mailbox server, which in most cases had sufficient resources to handle the load.

Exchange 2007 was released with dedicated roles for the various functions required in the Exchange organization. Among them was the CAS role, which as we know now, is a more mature version of the front-end server. The CAS role is intended to be the connection point for all client communications. Exchange 2007 accomplished this, for the most part. Outlook Web Access, Outlook Anywhere, POP, IMAP, Activesync, and Autodiscover connections all terminated with the CAS server from which connections were proxied back to the Mailbox server if required. The main difference now is where things such as content conversion, OWA page rendering, and client configuration were being processed. Those who thought of the CAS as

nothing more than an updated front-end server may have been a bit surprised when they realized that the CAS servers they deployed were under-provisioned and required either extra CPU, memory, or simply additional members to be added to the CAS pool.

The new CAS role was a great leap forward from the old front-end servers, but there was still one client connection point that was missing. MAPI connections from Outlook users within an organization's firewall were still terminating at the Mailbox servers. Exchange 2010 has solved this with RPC client access. Outlook clients connecting to an Exchange 2010 mailbox will now have to connect through the CAS server via RPC client access. This change improves the client experience by reducing the time it takes a client to reconnect after a failover. Instead of taking up to 15 minutes to reconnect, a client will now connect to their mailbox within 30 seconds of a failover.

With this new responsibility, you will need to make sure the CAS role is properly sized to accommodate the increase in load that will be placed by MAPI clients. In the following sections, you will be presented with the data required to size the CAS role appropriately and will be able monitor the role to find any potential bottlenecks.

Processor

With the CAS server taking on more responsibilities in the latest versions of Exchange Server, you will have to do some planning around processor configuration. In the default installation of Exchange, your CAS servers will be responsible for providing and rendering the Outlook Web Access (or Outlook Web App in Exchange 2010) interface, providing Autodiscover and Availability functions, and being the connection point for Outlook Anywhere (formerly RPC over HTTPS). If you are deploying in an Exchange 2010 environment, the default installation will include RPC Client Access. Finally, you still have the ability to provide Post Office Protocol (POP) and Internet Message Access Protocol (IMAP) capabilities to your end users, but this all requires more processing power.

Microsoft has again done a great job of laying down some initial baselines for Exchange administrators to use when planning their environment. With all the offloading of processing that has been done, the CAS role is able to efficiently use four virtual processors in environments where there is substantial non-MAPI client traffic. You also have the option to deploy in a scaled-out approach, which can provide the required amount of processing capabilities while providing redundancy. The recommended processor configuration for the CAS role is between four to eight virtual processors. The maximum recommended configuration is 12 virtual processors; this is currently more than is supported in a vSphere environment. The consideration on how many virtual processors to assign to a CAS role VM will come by looking at the recommended number of total CAS virtual processors and the number of VMs you desire to scale out to.

For Exchange 2007 and 2010, the recommended mailbox-to-CAS processor ratio is four-to-three. That is, for every four Mailbox role virtual processors, you should deploy three CAS role virtual processors. In an environment with two mailbox VMs, each with four virtual processors, you will require a total of six CAS role virtual processors. How you divide those virtual processors should depend on your scaling-out strategy. You can deploy one CAS role VM with six processors or three CAS role VMs with two virtual processors each. The ratio recommendation assumes all protocols are using Secure Sockets Layer (SSL) encryption and that the processor type and speeds are the same.

Memory

As with the increased processing requirement that the new roles have placed on the CAS, the memory usage has increased as well. The CAS role memory usage is mostly based on the number of client connections and the transaction rate at which those clients are processing data. This relationship between memory utilization and client connections is linear; the more client connections and transactions your CAS sees, the more memory you will utilize. As client connections drop, so will the utilized memory. The memory utilization and processor utilization are balanced as well, which means you will see your CAS become processor bound at the same time it becomes memory bound. Table 6.5 lists the current minimum, recommended, and maximum supported memory configurations for CAS servers.

Table 6.5: Memory Configurations for Exchange 2007/2010

Role/Version	Minimum Supported	Recommended	Maximum per Server
Client Access/ Exchange 2007	2 GB	1 GB per processor (2 GB minimum)	16 GB
Client Access/ Exchange 2010	4 GB	2 GB per core (8 GB minimum)	24 GB

The minimum supported memory configuration must be taken into account when designing your environment. As stated, it is the minimum supported configuration, and it is possible that support will be refused because of a lack of sufficient resources. On the contrary, the maximum per server is a recommendation based on price and performance and does not have any bearing on whether your design is supported.

Storage

Because of the mostly stateless nature of the CAS server, disk I/O is not a big concern for this role. As you have read in the previous couple of sections, the CAS role is pretty CPU and memory bound. That's not to say you shouldn't give disk I/O some thought. The basics still hold true; you want to ensure that the OS and page file have the resources they require to avoid a bottleneck. Consider some of the other processes that factor into disk I/O on a Client Access server:

Protocol logging The CAS acts as a web server, a POP server, and an IMAP server, and it provides various other services such as ActiveSync. With all these services, it is recommended that you enable and retain some level of logging to assist in troubleshooting or just to verify that the services are working as expected. This logging creates sequential write activity on the storage and consumes space.

Content conversion If any of your users are IMAP or POP clients, there's a good chance content conversion will come into play in your environment. When messages are sent through the transport engine, they are converted to MAPI and delivered to the mailbox store. When requested by an IMAP or POP client, the mail must be converted to Multipurpose

Internet Mail Extensions (MIME) before being sent back to the client. The CAS is responsible for this conversion. If the message is larger than 64 KB, the conversion takes place on disk, thereby adding I/Os to the storage.

The amount of additional disk I/O placed on the system by logging is minimal, but you should be aware of it. When testing your prototype VMs running the CAS role, be sure to have all the services and logging that you will run in production enabled so that you have a clear indication as to the load presented by your configuration.

Hub Transport

The workload characteristics of the Hub Transport role closely match those of the Edge Transport. Both roles are responsible for sending and receiving mail and applying policies to those messages such as antispam and antivirus filtering. The Hub Transport performs additional functionality that requires consideration when sizing for your environment. The main focus here will be to point out any additional workloads present in the Hub Transport role that are not present in the Edge Transport role. This section, coupled with the Edge Transport section, will provide the knowledge required to deploy transport services that can help meet your organization's requirements.

EdgeSync

The EdgeSync process uses the pairing of the Edge Transport and Hub Transport services to synchronize directory data. This directory data contains the mail-enabled objects (recipient and distribution lists) as well as configuration data for the Exchange organization. The transfer of data happens at the Hub Transport servers and is pushed out to the Edge Transport servers. To facilitate this transfer, the data is held in memory on the Hub Transport servers. Approximately 4 KB of memory per mail-enabled object is consumed by the EdgeSync process. When sizing the memory for the Hub Transport role, consider the additional memory required based on the number of mail objects that are present in the directory as well as growth expectations.

Transport Dumpster

The transport dumpster was introduced with Exchange 2007 to support the new continuous replication features. This feature has carried over to Exchange 2010 for use with protected databases (those in database availability groups). Because of the asynchronous nature of continuous replication, there is the chance that some log files will be missed in the event of a lossy failover. To mitigate the risk of lost mail during a lossy failover, the transport dumpster retains a copy of each message where at least one of the recipient's mailboxes is stored on a protected database (LCR, CCR, DAG).

The I/O footprint of the Hub Transport with the transport dumpster enabled is significantly increased. The storage that holds the queue database will not only have an increase in write activity, but reads will also occur. In Table 6.6, an average message size of 40 KB is used to show the difference in I/Os per message with and without the transport dumpster enabled. Although originally published for Exchange 2007, these numbers provide a good baseline for an Exchange 2010 design.

TABLE 6.6: Hub Transport I/Os per Message

HUB TRANSPORT SERVER DATABASE I/O (STEADY STATE)	TRANSPORT DUMPSTER ENABLED	TRANSPORT DUMPSTER DISABLED
Total IOPS per message (approximately 40 KB)	17	4
Log write I/Os per message (sequential)	7	2
Database write I/Os per message (random)	7	2
Database read I/Os per message (random)	3	0

Microsoft TechNet: `http://technet.microsoft.com/en-us/library/bb738141(EXCHG.80).aspx`

With the numbers from Table 6.6, you can use the same formula that was used for the Edge Transport role to calculate the database and transactional I/O requirements for both scenarios. Plug in the desired message rate, such as 20 messages per second, plus 20 percent for growth and traffic spikes.

Transport Dumpster Enabled

DB IOPS per message (TD enabled) = (20 × (7 + 3)) + 20% = 240

LOG IOPS per message (TD enabled) = (20 × 7) + 20% = 168

Transport Dumpster Disabled

DB IOPS per message (TD enabled) = (20 × (2 + 0)) + 20% = 48

LOG IOPS per message (TD enabled) = (20 × 2) + 20% = 48

You can see the difference enabling the transport dumpster makes on the overall I/O requirements. In reality, this usually isn't something that sways a decision to use continuous replication. The fact is that the extra load placed by the transport dumpster will usually be easily handled by a redundantly configured storage subsystem with multiple disk spindles. Regardless, it is important to keep this data handy when designing your solution to ensure your Hub Transport can satisfy the requirements.

The transport dumpster queue is part of the Hub Transport queue database (`mail.queue`). The size of the transport dumpster is an organization-wide configuration parameter. To view the current configuration, use the Exchange Management Shell command `get-transportconfig`.

```
[PS] C:\>get-transportconfig


ClearCategories                : True
DSNConversionMode              : UseExchangeDSNs
GenerateCopyOfDSNFor           : {5.4.8, 5.4.6, 5.4.4, 5.2.4, 5.2.0, 5.1.4}
InternalSMTPServers            : {}
JournalingReportNdrTo          : <>
```

```
MaxDumpsterSizePerStorageGroup : 18MB
MaxDumpsterTime                : 7.00:00:00
MaxReceiveSize                 : 10MB
MaxRecipientEnvelopeLimit      : 5000
MaxSendSize                    : 10MB
TLSReceiveDomainSecureList     : {}
TLSSendDomainSecureList        : {}
VerifySecureSubmitEnabled      : False
VoicemailJournalingEnabled     : True
HeaderPromotionModeSetting     : NoCreate
WritingBrandingInDSNEnabled    : True
Xexch50Enabled                 : True
```

By default the Exchange 2007 transport dumpster is enabled and configured for 18 MB per protected storage group with a seven-day retention period. The retention period is the amount of time that messages are kept in the transport dumpster. Seven days is the recommended setting in Exchange 2007. Exchange 2010 introduces a new way of managing the size of the transport dumpster. Instead of keeping messages for a set period of time, messages are tracked via the replication pipeline to determine which messages have been delivered and replicated. Once a message has been confirmed delivered and replicated, the message is truncated from the transport dumpster.

The transport dumpster should be sized to your environment. The recommended setting for the transport dumpster is 1.5 times the maximum allowed message size. In other words, if your maximum message size limit is 10 MB, it is recommended to set the transport dumpster size to 15 MB. In environments where there are no message size restrictions, the recommendation is to size the transport dumpster to 1.5 times the average message size.

Each Hub Transport server in the Active Directory site will maintain the globally configured transport dumpster size per protected storage group. In other words, if your environment has 10 protected storage groups, 2 Hub Transport servers, and a transport dumpster size of 10 MB, each Hub Transport will dedicate 100 MB of queue database space to the transport dumpster (10 storage groups × 10 MB). This space must be accounted for when sizing the storage for the queue database. Additionally, each storage group will have an aggregate amount of transport dumpster space spread across all Hub Transport servers. In the previous example, each storage group would have an aggregate of 20 MB of transport dumpster space (10 MB of transport dumpster per storage group × 2 Hub Transport servers).

To set the transport dumpster to suit your environment, use the `set-transportconfig` Exchange Management Shell command:

```
[PS] C:\>Set-TransportConfig -MaxDumpsterSizePerStorageGroup 15MB -MaxDumpsterTi
me 14.00:00:00

[PS] C:\>get-transportconfig


ClearCategories                : True
DSNConversionMode              : UseExchangeDSNs
GenerateCopyOfDSNFor           : {5.4.8, 5.4.6, 5.4.4, 5.2.4, 5.2.0, 5.1.4}
InternalSMTPServers            : {}
```

```
JournalingReportNdrTo          : <>
MaxDumpsterSizePerStorageGroup : 15MB
MaxDumpsterTime                : 14.00:00:00
MaxReceiveSize                 : 10MB
MaxRecipientEnvelopeLimit      : 5000
MaxSendSize                    : 10MB
TLSReceiveDomainSecureList     : {}
TLSSendDomainSecureList        : {}
VerifySecureSubmitEnabled      : False
VoicemailJournalingEnabled     : True
HeaderPromotionModeSetting     : NoCreate
WritingBrandingInDSNEnabled    : True
Xexch50Enabled                 : True
```

Client Access and Hub Transport Combined Role

In many organizations, it may not make sense to have separate Hub Transport and Client Access servers. This is especially true in organizations that are going through a server consolidation exercise, have existing underutilized Hub Transport or Client Access servers, or that are virtualizing these roles. Smaller organizations can provide redundancy and consolidate by combining these roles in a virtual environment. Combined roles have different sizing requirements that are explained in the following sections.

Processor

If you refer to the guidance for stand-alone Client Access or Hub Transport servers, you can see that the processor sizing for these roles is mostly calculated as a ratio to the active mailbox processors in the site. When combining these roles, the ratio becomes a simple one-to-one ratio, that is, one CAS/Hub Transport combined role processor to one Mailbox role processor. This makes sizing combined roles trivial. A small organization may choose to deploy a 4-vCPU Mailbox role VM to support 2,000 average users and deploy two 2-vCPU combined role VMs to satisfy the processor ratio and provide CAS and Hub Transport redundancy.

As with other roles, Microsoft has minimum, recommended, and maximum processor requirements and recommendations. If you choose to deploy a combined role virtual machine, keep in mind that the minimum requirement to receive support is two vCPUs. As with the Mailbox role, the recommended configuration is four vCPUs. The maximum recommended configuration is 12 vCPUs. Deploying on more than 12 vCPUs is not recommended.

Memory

The combined roles do not present any significant changes in the required memory. The general guidance is to configure your dual-role VM with at least 4 GB of memory; this is the minimum requirement. The recommended memory configuration for a virtualized dual-role virtual machine is 8 GB. This aligns with the recommended four-vCPU configuration. The recommended maximum for a combined role virtual machine is 2 GB per vCPU allocated to the virtual machine. If deployed at the maximum size allowed in vSphere, a combined-role virtual machine would have eight vCPUs and 16 GB of memory.

If choosing to deploy combined role servers, keep in mind that if your current environment experiences high usage in the CAS role or Hub Transport role, that spike could negatively affect the other roles on your system. A good example of this is the IMAP service on the Client Access role. When fetching a large folder via IMAP, the MSExchangeIMAP4 service can consume 100 percent of the available CPU resources to process the large fetch request. This can impact the Hub Transport role and cause a queue to build while the IMAP request is processed. If working on a new deployment, it is recommended that you deploy separate role virtual machines and attempt to combine only after realizing the baseline load placed on your systems during normal work hours.

Mailbox Server

Client perception of the messaging environment you support will likely come down to the performance of the Mailbox server role. Operations performed by the end users will trigger activity to the back-end Mailbox server. For this reason, the Mailbox server is considerably the most important role to size correctly.

Exchange 2010 was designed with architectural changes that significantly improved performance. Building on the improvements made in Exchange 2007 such as the 64-bit design, Exchange 2010 continues to make great strides. Database I/Os are now larger and more sequential, and the database page size has been increased from 8 KB to 32 KB. The larger page size reduces disk I/O and improves performance by caching the larger pages in memory.

Sizing the Mailbox server can be tricky, but Microsoft provides great guidance for processor, memory, and storage design. The Exchange team has kept the Exchange Server Role Calculator updated to make this process even easier. By knowing a few key metrics from an existing Exchange environment or using some baseline numbers built into the calculator, you can quickly have the requirements in hand to assist in further planning. When deploying in a virtual environment, the same compute resource requirements exist. Using the sizing guidance, you will be able to properly size your virtual Mailbox servers to provide the expected performance for your environment.

Processor

Exchange 2010 introduces some new concepts around high availability that require you to put a bit more effort into the processor design of your Mailbox server role. DAGs, along with the concept of supporting both active and passive database copies on the same Exchange 2010 Mailbox server, mean you must take the additional processing required by passive database copies into consideration. Not only is the Mailbox server role responsible for the back-end processing of end-user mailbox data, but a Mailbox server supporting DAGs must also perform the validation and replay of log files as well as maintain the content index for its active and passive database copies. For every additional copy of a database, the per-mailbox CPU requirement must be increased by 10 percent on the host serving the active mailbox database.

Table 6.7 gives an estimate of the database cache, IOPS, and CPU requirements based on the user profile for Exchange 2010. In this case, the user profile is based on messages sent and received per day. CPU requirements in Table 6.7 are listed in megacycles per active and passive mailbox. The general recommendation is that a passive mailbox will require approximately 15 percent of the megacycles as the active mailbox.

TABLE 6.7: Estimated Database Cache, IOPS, and CPU Requirements by User Profile

Messages Sent or Received per Mailbox per Day	Database Cache per Mailbox in Megabytes (MB)	Single Database Copy (Stand-Alone) with Estimated IOPS per Mailbox	Multiple Database Copies (Mailbox Resiliency) with Estimated IOPS per Mailbox	Megacycles for Active Mailbox or Stand-Alone Mailbox	Megacycles for Passive Mailbox
50	3	0.06	0.05	1	0.15
100	6	0.12	0.1	2	0.3
150	9	0.18	0.15	3	0.45
200	12	0.24	0.2	4	0.6
250	15	0.3	0.25	5	0.75
300	18	0.36	0.3	6	0.9
350	21	0.42	0.35	7	1.05
400	24	0.48	0.4	8	1.2
450	27	0.54	0.45	9	1.35
500	30	0.6	0.5	10	1.5

Microsoft TechNet: `http://technet.microsoft.com/en-us/library/ee712771.aspx`

Megacycles per Mailbox

Table 6.7 is based on the Intel Xeon x5470 processors in a 2 × 4 core configuration. These processors have a clock speed of 3.33 GHz per processor core and deliver about 3,300 megacyles of performance throughput. For calculating the required megacycles when using any other type of processor, a megacycle adjustment is required. The adjusted value is then entered into the Exchange 2010 Server Role Calculator. To calculate the adjusted value, follow these steps:

1. Browse to `www.spec.org`, and click Results.
2. On the Published SPEC Benchmark Results page, locate CPU2006, and click the Simple CPU2006 Result Search link.
3. In the Available Configurations field, select SPECint2006 Rates, and select the Simple radio button.
4. In the Simple request section, select Processor from the drop-down of search criteria, enter the processor type (such as **x5550**), and click Search.

5. In the results, find the SPECint_rate2006 section, and locate the server and processor you will be using, such as Dell R710 with the Intel x5550 2.67GHz processor.
6. Make note of the result (252 or 31.5 per core), the number of cores (8), the number of chips (2), and the number cores per chip (4).

   The baseline system (HP DL380 G5 x5470 3.33 GHz × 8 cores) has a SPECint_rate2006 results value of 150, or 18.75 per core.
7. To determine the megacycles of the R710 platform, use the following formula:

```
Megacycles per core = ((new platform per core value) * (total megacycles of new
platform per core)) / (baseline per core value)
(31.5 * 2670) / 18.75 = 4485.6 megacycles per core
```

8. This number is then plugged into step 5 of the Exchange 2010 Exchange Role Calculator.

By using this formula, you can determine exactly how many processor cores you will require when sizing your Exchange environment. Using the default of 3,300 megacycles will most likely yield a conservative result since most enterprise-level processors will be able to produce a higher megacycle throughput. If you are planning on deploying Exchange 2007 in your environment, the process is slightly simpler.

In Exchange 2007, the concept of Database Availability Groups does not exist, nor does the concept of replication content index data. This significantly reduces the per-mailbox CPU requirements. Additionally, Exchange 2007 still supports the shared-storage model for high availability known as *single-copy clusters*. The rule of thumb for sizing CPU requirements in Exchange 2007 revolved mainly on mailbox profile, mailbox count per server, and whether any third-party utilities were to be used such as Microsoft Forefront.

From a budgeting perspective, the guidance from Microsoft about mailboxes per core for Exchange 2007 has been 1,000 average users per one CPU core. This model allows for up to 4,000 users on a 4× core server. In a virtual environment, these numbers work out quite well given the size of physical server hardware that is available these days. With a maximum recommendation of 12 CPU cores per Mailbox server, companies that have standardized on higher-end physical servers containing 16 or more CPU cores either have to adjust their standards or have to over-commit hardware. Partitioning the larger physical servers into smaller virtual machine units allows these companies to efficiently use the resources they have purchased.

Table 6.8 provides the current minimum supported recommended, and maximum recommended processor configuration for the Exchange Mailbox role for both Exchange 2007 and 2010. The maximum recommended is not a support consideration, but rather a balance between cost and performance.

TABLE 6.8: Processor Configurations for Exchange Mailbox roles

SERVER ROLE	MINIMUM	RECOMMENDED	MAXIMUM RECOMMENDED
2010 Mailbox	2 × processor core	4× processor cores	12× processor cores
2007 Mailbox	1× processor core	4× processor cores	12× processor cores

Memory

The guidance about providing memory for Exchange 2010 has changed quite significantly from that given for Exchange 2007. Not only has the maximum usable memory cache per mailbox increased (from 7 MB to 30 MB), but the per-database cache has been increased from 20 MB to 100 MB for databases with multiple copies. This increase in database cache directly translates to reduced disk I/O by storing more database changes in memory, thus providing improved coalescence of disk I/O, which equals overall fewer disk I/Os.

To make calculating the amount of usable mailbox cache available while providing the required memory for the OS and applications, refer to Table 6.8 and Table 6.9. As was the case in previous versions of Exchange, the number of databases (previous versions referred to storage groups) determines the maximum amount of memory required. Table 6.9 provides the minimum memory requirements given the number of databases (active and passive).

Table 6.9: Minimum Memory Required per Database (Active and Passive)

Database Count	Exchange 2010 Minimum Required Physical Memory
1–10	2 GB
11–20	4 GB
21–30	6 GB
31–40	8 GB
41–50	10 GB
51–60	12 GB
61–70	14 GB
71–80	16 GB
81–90	18 GB
91–100	20 GB

Chances are you will be serving more mailboxes than the minimum requirements listed in the table could handle, but keep the numbers in mind because they are the minimum requirements. Table 6.10 provides the recommended database cache based on the user profile (messages sent and received per day) and the estimated IOPS for a stand-alone database and multiple-copy databases. Note the difference in IOPS between single and multiple databases. This is because of the increased log checkpoint depth target (100 MB) allowed for multicopy databases.

TABLE 6.10: Estimated IOPS Based on Mailbox Profile and Mailbox Database Cache

Messages Sent/ Received per Mailbox per Day (~75KB Average Message Size)	Database Cache per User (MB)	Single Database Copy (Stand-Alone): Estimated IOPS per Mailbox	Multiple Database Copies (Mailbox Resiliency): Estimated IOPS per Mailbox
50	3	0.06	0.05
100	6	0.12	0.1
150	9	0.18	0.15
200	12	0.24	0.2
250	15	0.3	0.25
300	18	0.36	0.3
350	21	0.42	0.35
400	24	0.48	0.4
450	27	0.54	0.45
500	30	0.6	0.5

With the previous tables you can determine the minimum memory requirements based on the number of databases you will house on a single Mailbox server and the recommended database cache per mailbox based on user profile. The final piece of this memory puzzle is to provide adequate memory to the Mailbox server to ensure the required mailbox database cache is satisfied as well as the minimum server requirements. Table 6.11 lists the database cache size provided to the Mailbox server role given the amount of physical memory installed. Both mailbox-only and multirole (Mailbox + Hub Transport) servers are covered in this list.

TABLE 6.11: Mailbox Database Cache Sizes

Server Physical Memory (RAM)	Database Cache Size (Mailbox Role Only)	Database Cache Size Multiple-Role (for example, Mailbox + Hub Transport)
2 GB	512 MB	Not supported
4 GB	1 GB	Not supported
8 GB	3.6 GB	2 GB
16 GB	10.4 GB	8 GB

TABLE 6.11: Mailbox Database Cache Sizes *(CONTINUED)*

SERVER PHYSICAL MEMORY (RAM)	DATABASE CACHE SIZE (MAILBOX ROLE ONLY)	DATABASE CACHE SIZE MULTIPLE-ROLE (FOR EXAMPLE, MAILBOX + HUB TRANSPORT)
24 GB	17.6 GB	14 GB
32 GB	24.4 GB	20 GB
48 GB	39.2 GB	32 GB
64 GB	53.6 GB	44 GB
96 GB	82.4 GB	68 GB
128 GB	111.2 GB	92 GB

STORAGE

The latest versions of Exchange have reduced disk I/O by as much as 90 percent compared to Exchange 2003 and previous versions. The release of 64-bit only versions of Exchange allowed for a greatly increased mailbox database cache. As described in the previous section, increased mailbox database cache directly affects disk I/O by reducing the need to go to disk for frequently accessed database pages. Database page size was increased from 4 KB in Exchange 2003 and older to 8 KB in 2007 and now to 32 KB in Exchange 2010; again, all of this helps achieve that reduction in I/O, which helps in designing your storage environment.

Regardless of how much the I/O required per mailbox is reduced, there are still latency concerns, which unless addressed will impact the end user's experience. User profile, mailbox size, third-party utilities, and mobile devices all contribute to the I/O footprint presented to your Mailbox servers. Although Microsoft has done a great job at reducing how many disk spindles are required to provide acceptable disk latency, you don't just want to ask the storage administrators for a few hundred gigabytes of space and be on your way. This is especially true in the virtual environment where there is a tendency to mask the underlying physical resources.

Mailbox Role Performance Characteristics

Understanding what contributes to the I/O generated on the Mailbox role is key to designing the proper storage configuration. Not only do you have to take into consideration the load placed by your end users, but now you have the ability to have multiple copies (up to 16) of a single database by using DAGs. The following list provides the factors you must consider when designing the underlying storage:

Mailbox count The number of users directly impacts the I/O per server and capacity requirements.

Mailbox size Setting a mailbox quota will greatly improve the predictability of required storage and future growth based on mailbox count.

Concurrency Knowing the concurrent number of users will help in determining how much active I/O to plan for. If the server hosts users from differing geographic locations, you may be able to leverage this to more efficiently utilize disk performance.

Mailbox profile The mailbox profile is determined by the amount of mail sent and received per user per day and the average message size. The profile type directly impacts transactional I/O.

Outlook mode Outlook users have the option to used online or cached mode to connect to their mailboxes. Cached mode significantly reduces the I/O load placed on the Mailbox server and allows the users to work offline.

Third-party applications Tools such as client-side search tools may index the online copy of a mailbox or public folders, thus increasing load on the server. Utilities that run against the server such as third-party mobile device applications or antivirus software can also lead to increased I/O.

Data replication Whether the replication is storage based or you choose to use a built-in mechanism such as CCR or DAG, there will be an increase in I/O at the server level.

You can collect this information in numerous ways. If you have an existing environment, you can use tools such as the built-in performance monitor in Windows Server 2008 to gather the counters you want to measure and review the results manually. One tool that simplifies this process is the Microsoft Exchange Server Profile Analyzer. The Profile Analyzer tool collects statistical information from your production mailbox stores and can provide information such as mailbox size, rule count, messages per mailbox, and average message size. For more information, refer to this Microsoft TechNet article:

```
http://technet.microsoft.com/en-us/library/bb508856(EXCHG.65).aspx
```

For new environments or environments transitioning from a non-Exchange solution or a platform where the numbers don't exactly translate into usable information, the Exchange team has provided a tool that can be used to size not only the storage but also most aspects of your Exchange Mailbox role. The Exchange 2010 Server Role Requirements Calculator (Figure 6.5) has been the de facto sizing tool used by most Exchange implementers for years. It is developed and maintained by the Microsoft Exchange Team, and who would know best how the Mailbox role should be deployed? The calculator has been updated for Exchange 2010 and should be used even if you have your own performance metrics, at least to validate your design against the recommended best practices.

You can find the Exchange requirements calculator at this URL:

```
http://go.microsoft.com/fwlink/?LinkId=178613
```

When choosing to deploy a new Exchange environment using replicated databases, whether it be CCR or DAG, it is important to consider the impact replication has on storage. The Exchange Server Role Requirements Calculator takes into account the additional I/O and processing overhead that continuous replication adds to the Mailbox server role. With this data in hand, you and your storage vendor can design the underlying storage design that suits your sizing and performance requirements. Fortunately, virtualization has not changed the design aspects of the underlying storage. How we present and attach the storage is affected by virtualization, and we'll touch on the relevant aspects of this in the next section.

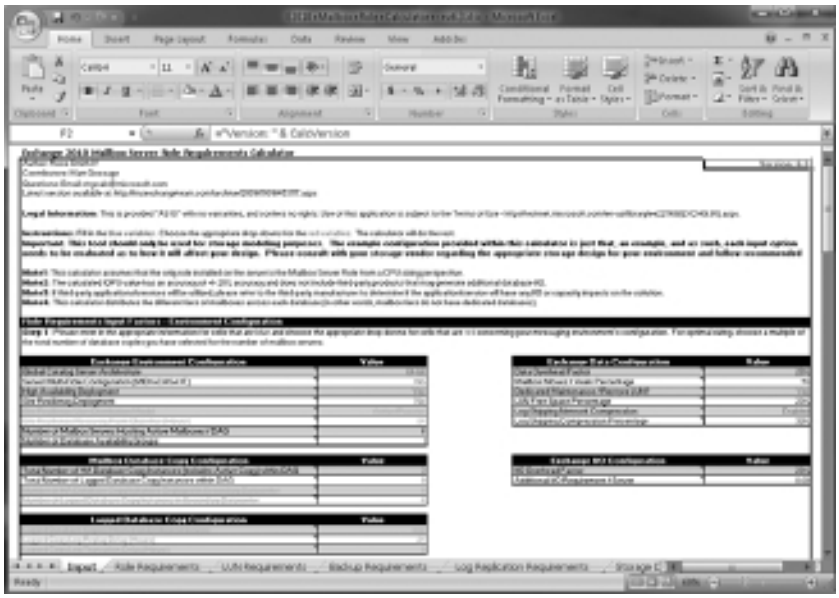

FIGURE 6.5
Exchange 2010 Server Role Requirements Calculator

VMFS vs. RDM

The topic of VMFS partitions and raw device mappings should not be new. Chapter 3 discussed the two briefly in respect to deploying Microsoft Cluster Services and failover clustering in a virtual environment. You'll recall when deploying a shared storage cluster in a virtual environment, the recommendation for enterprise environments was to deploy a cluster using virtual machines across ESX hosts. This required the use of physical mode raw device mappings (RDMs). Clusters aren't the only reason you need to consider deploying raw device mappings.

A common misconception is that RDMs, also described as pass-through storage devices (because the path from the virtual machine to the storage passes through the hypervisor with no virtualization interaction), perform better than a virtual disk file residing on a VMFS. Although this may have been the case in earlier versions of many hypervisors, the VMFS file system has been tuned to the point where latency times between VMFS and RDMs are within microseconds of each other. These days, performance rarely plays a role when deciding whether to deploy on VMFS or RDM.

Deciding on the storage access method to use should be based on your design requirements. Common factors include support, such as in the case of clustering. To deploy a Microsoft cluster, you must use physical-mode RDMs as the shared disks. How you plan to back up the data, in this case the Exchange databases and logs, can affect the method of storage access. Many Exchange deployments rely on the VSS snapshot technology introduced with Windows 2003. Hardware-based VSS solutions from the leading storage vendors often require that a backup agent be installed on the Exchange Mailbox server to initiate the VSS activities. These agents require direct access to the back-end storage devices to achieve the cloning or snapshotting at the storage layer once Exchange has been quiesced. This type of access will, in most cases, rely on physical-mode RDMs.

Raw device mappings come at a price. Because the virtualization layer is removed (in physical-mode RDMs), the ability to clone, migrate, and initiate a VMware Data Recovery backup and then snapshot the entire VM is lost. Migration via VMotions will continue to work, but many of the other features available by having the entire virtual machine being comprised of a set of files are lost.

Although the Exchange mailbox database has been optimized to reduce the load on the underlying storage, there is still plenty to consider when designing a new environment. Fortunately for you, virtualization does not add much to the already complex process. Understanding the storage technical requirements up front, knowing the support considerations, and working with your storage vendors can significantly ease the burden of this cornerstone in your virtualized Exchange design.

Unified Messaging

The Unified Messaging role (UM) was first introduced with Exchange 2007 and has been carried into Exchange 2010. UM allows for the integration of your company's telephony (voice and fax) and email system. UM users are able to seamlessly access their voicemail, access their email, and send and receive faxes using a single client. Unified Messaging servers should be placed in close proximity to the PBX system that they will integrate with and optimally on the same VLAN or broadcast domain as the IP gateway, which will perform the translation to Voice over IP (VoIP). Unless quality of service has been configured on the local-area network to ensure the quality of voice calls traversing the data network, it is possible your organization may employ multiple networks dedicated to either voice or data. In these scenarios, virtualizing the UM role of Exchange may require dedicated networking equipment (both physical and virtual) to avoid having to go through multiple network hops to reach the VoIP network. In a virtual environment, this may look something like Figure 6.6. The virtual machine acting as the UM role has two virtual network adapters, one accessing the data network and the other accessing the VoIP network.

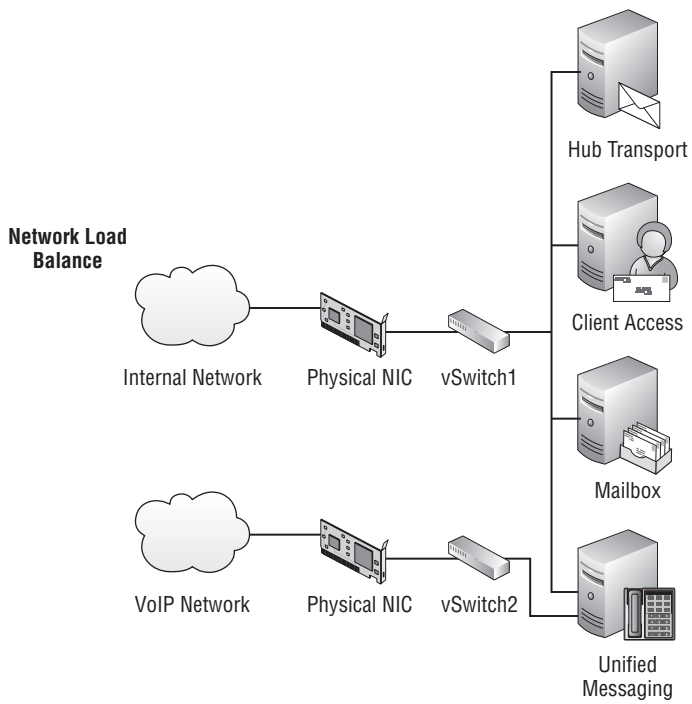

FIGURE 6.6
Dual-homed UM virtual machine accessing both the data and VoIP networks

Unfortunately, Microsoft has chosen not to support the Unified Messaging role in a virtual environment. Because of this, there has not been much adoption to the concept of virtualizing this role. In the following sections, we have listed the resource recommendations that will directly translate to a virtual environment; however, because of the lack of support, we recommend not virtualizing this role until Microsoft sees it as a supportable configuration.

Processor

Table 6.12 lists the minimum and maximum processor recommendations for the Unified Messaging role.

Table 6.12: Processor Recommendations for Unified Messaging

Server Role	Minimum Supported	Maximum Recommended
Exchange 2010 Unified Messaging	2× processor cores	12× processor cores
Exchange 2007 Unified Messaging	1× processor core	6× processor cores

Memory

Table 6.13 lists the minimum and maximum memory recommendations for the Unified Messaging role.

Table 6.13: Memory Recommendations for Unified Messaging

Server Role	Minimum Supported	Recommended
Exchange 2010 Unified Messaging	4 GB	2 GB per core
Exchange 2007 Unified Messaging	2 GB	1 GB per core

Storage

The Unified Messaging role, much like the Client Access server, is mostly stateless and requires a fraction of the disk I/O required by the Hub Transport or Mailbox roles. Because voicemails and faxes are kept in the end users' mailboxes, the capacity requirements for the UM servers are minimized. The source of disk I/O for the UM role will mostly come from content conversion and protocol logging that has been enabled. Incoming voicemails are converted from the native recorded format to a compressed format defined by the administrator such as MP3. This conversion requires more processing resources than any other, but having well-responding storage will ensure the conversions are as efficient as possible.

Although not being formally supported in a virtualized deployment, virtualizing the UM role in a test environment will provide for a cost-effective development, learning, and testing solution. Testing configuration updates and OS and application updates in a test bubble before deploying to production is usually a requirement for most organizations' change control process. Having an easy-to-replicate and rollback test environment has proved invaluable for many organizations to achieve this process requirement.

Monitoring Performance

During deployment and testing, you will want to verify that the system is running as efficiently as possible. Using tools like Service Center Operations Manager can help aid in monitoring, but for real-time analysis or more precise figure capturing, you will want to take advantage of the built-in performance counters. In Tables 6.14–6.17, you will find the counters you can focus on to determine whether your Exchange server is running as efficiently as possible, whether there is a bottleneck in the system, and, if so, which component may be the culprit (according to Microsoft TechNet article `http://technet.microsoft.com/en-us/library/bb201720(EXCHG.80).aspx`). These include processor, memory, storage, and network counters.

Table 6.14: Processor Counters

Counter	Expected Values
Processor(_Total)\% Processor Time Shows the percentage of time that the processor is executing application or operating system processes.	Should be less than 75 percent on average

Table 6.15: Memory Counters

Counter	Expected Values
Memory\Available Mbytes Shows the amount of physical memory, in megabytes (MB), immediately available for allocation to a process or for system use.	Should remain above 100 MB at all times
Memory\Pages/Sec Shows the rate at which pages are read from or written to disk to resolve hard page faults.	Should be below 1,000 on average

TABLE 6.16: Storage Counters

COUNTER	EXPECTED VALUES
Logical/PhysicalDisk(*)\Avg. Disk sec/Read	Should be less than 20 milliseconds (ms) on average.
Shows the average time, in seconds, it takes to read data from the hard disk.	Spikes (maximum values) should not be higher than 50 ms.
Logical/PhysicalDisk(*)\Avg. Disk sec/Write	Should be less than 20 milliseconds (ms) on average.
Shows the average time, in seconds, of a write of data to the disk.	Spikes (maximum values) should not be higher than 50 ms.
Logical/Physical Disk(*)\Avg. Disk sec/Transfer	For healthy hard disks, this counter shows approximately 20 ms. Counter values larger than 20 ms, or with large spikes, indicate a possible hard disk issue (for example, failure or slow speed).
Indicates how fast data is being moved (in seconds). Measures the average time of each data transfer, regardless of the number of bytes read or written.	

TABLE 6.17: Network Counters

COUNTER	EXPECTED VALUES
Network Interface(*)\Bytes Total/sec	For a 100-Mbps network adapter, should be below 6–7 Mbps.
Indicates the rate at which the network adapter is processing data bytes.	For a 1000-Mbps network adapter, should be below 60–70 Mbps.
This counter includes all application and file data, in addition to protocol information such as packet headers.	
Network Interface(*)\Packets Outbound Errors	Should be 0 at all times.
Indicates the number of outbound packets that could not be transmitted because of errors.	

The general counters in the previous tables will provide a baseline of understanding how your Exchange servers are performing at the basic compute level. The list of counters for each role can be quite long, and Microsoft has done a great job of publishing these on its TechNet website at the following URL: `http://technet.microsoft.com/en-us/library/dd335215.aspx`.

Gathering all the data provided by these counters can be a daunting task, not to mention actually taking that data and trying to analyze the results. Luckily, there is a great (and free) tool that you can use to do the heavy lifting and provide some great insight into the performance of your Exchange roles.

Performance Analysis of Logs Tool

The Performance Analysis of Logs (PAL) tool uses VBScript and the Microsoft Log Parser tool (another free download) to input previously collected performance logs and analyze the results against known thresholds. The tool will calculate the results and display them in a simple-to-digest HTML report. The PAL tool provides analysis for most Microsoft tier-1 applications.

To use PAL to analyze your environment, you will require the following components:

- Log Parser 2.2: `www.microsoft.com/downloads/details.aspx?FamilyID=890cd06b-abf8-4c25-91b2-f8d975cf8c07&DisplayLang=en`
- PAL Tool: `www.codeplex.com/PAL`
- Perfwiz Replacement XML files: `http://blogs.technet.com/mikelag/archive/2008/05/02/perfwiz-replacement-for-exchange-2007.aspx`

The high-level overview of the process to create a PAL report is as follows:

1. Download and install Log Parser and the PAL tool.
2. Install the Perfwiz replacements, either the role based or full counters.
3. Run the performance counters during peak times or while troubleshooting a performance issue.
4. Launch the PAL wizard.
5. Create the PAL job and execute.
6. Review the HTML file created by the PAL tool.

The PAL tool is a valuable resource to add to your toolkit for quick analysis of your Exchange environment. Using the PAL tool on a regular basis to provide historical baselines along with proactive monitoring with tools such as Microsoft System Center Operations Manager will prove useful when troubleshooting any performance problems in your environment.

Backup and Restore

When looking to protect Exchange data, administrators will focus on the Mailbox role. After all, this is where the majority of your data will reside. Depending on the size of the organization, this may involve backing up tens of terabytes (TB) worth of Exchange databases. When dealing with this much data, it is important to implement a solution that will provide a consistent backup and restore method. This will include being able to complete backup operations without interfering with normal user operations.

As mentioned in prior chapters, virtualizing your environment will provide additional methods for protecting your data, but that doesn't mean that traditional backup methods cannot be used. In fact, and especially for Exchange, the traditional agent-based backup methods continue to be the number-one method for backing up Exchange databases. The reason for this is that as

an Exchange administrator you understand that the Exchange-based backup process provides additional layers of protection and maintenance. Not only does a successful backup truncate the log files that are no longer required, but it also provides an integrity check that if found can lead to early warning signs of issues to come. For this reason, the recommended backup method for the Exchange Mailbox role continues to be those that employ the Exchange backup APIs, whether they use the streaming backup method (Exchange 2007) or are VSS based (Exchange 2007 to 2010). Most storage vendors provide hardware-based VSS solutions that will integrate with tier-1 Microsoft applications. Microsoft and other third-party data protection vendors can provide software-based VSS solutions. Regardless of the method your organization chooses to deploy, you can rest assured that virtualization will not impede your ability to protect your company's most valuable assets.

VMware Exchange Building Blocks

One of the most daunting tasks for any Exchange/systems administrator is sizing Exchange for performance and capacity. They will oversize their environments to ensure that their current demand and capacity needs are met with some percentage of headroom built-in.

One of major goals in the creation of the VMware Exchange building blocks was to simplify the sizing of Exchange servers. We took the Microsoft best practices, created modular virtual machines, and performed internal testing to ensure they would perform as Microsoft said they would.

The best practices guidelines come from the Microsoft TechNet document "Planning Your Server and Storage Architecture." After our initial internal testing, we set out to validate our findings with our various server and storage vendors.

TIP For more information on planning for Exchange Server, please the Microsoft TechNet site at the following URLs:

For Exchange Server 2010:

```
http://technet.microsoft.com/en-us/library/aa998636.aspx
```

For Exchange Server 2007:

```
http://technet.microsoft.com/en-us/library/aa998636(EXCH.80).aspx
```

As a VMware best practice for any customer looking to virtualize Exchange, we recommend using the building-block approach because of its modular and flexibility. Using pre-sized Exchange virtual machines, you can choose to scale up or out using one or many virtual machines on a single ESX server host all while achieving very predictable performance.

VMware Exchange building blocks have been thoroughly tested by all server and storage manufacturers, and all forms of storage were tested and validated to include the Network File System (NFS). Although NFS is no longer officially supported by Microsoft for Exchange storage, we wanted to gauge the performance of NFS in relation to iSCSI storage. We found NFS can actually perform as well as iSCSI and, in some cases, marginally better.

This whole idea of VMware Exchange building blocks was to create a modular approach for deploying Exchange. In testing with our partners, we ensured that the building blocks performed consistently across all storage and server vendors. If you follow the guidance for using the building blocks, you will achieve performance that rivals a physical server deployment of Exchange.

Through the use of VMware Exchange building blocks, you don't have to worry about placing all of your eggs in one basket because your virtual machines can be migrated to other hosts in the event of host outage, planned maintenance, and increased resource requirements. When a virtualized Mailbox server requires more memory or virtual CPU (vCPU), just live migrate it using VMotion to a more powerful ESX host.

TIP For more information on the VMware Exchange building blocks, please visit the VMware website:

`www.vmware.com/files/pdf/solutions/best_practices_exchange_solutions.pdf`

Listed in Table 6.18 are the Exchange building blocks for an average Exchange user.

TABLE 6.18: Building Block for CPU and RAM (Average Profile)

BUILDING BLOCK	500 USERS	1000 USERS	2000 USERS	4000 USERS
Profile	Average	Average	Average	Average
vCPU	1	1	2	4
Base RAM (GB)	2	2	2	2
Mailbox RAM(MB)	3.5 per mailbox	3.5 per mailbox	3.5 per mailbox	3.5 per mailbox
Total memory (GB)	4	6	9	16

These figures are for the average profile users. We actually followed the building-block approach when designing and implementing our internal deployment. Using this approach, you can work around the 8 core/32 GB recommended maximum for Exchange Server 2007 by "stacking" virtual machines on the same ESX host.

We have worked with partners to create test frameworks that allow us to showcase 16,000 "heavy" Exchange users on a single ESX host. We have worked with several customers to help them virtualize Exchange. For example, one customer is currently virtualizing 750,000 mailboxes on Virtual Infrastructure 3 and Exchange Server 2003. Other customers have successfully virtualized Exchange 2007 and now 2010 using the VMware Exchange building blocks.

Be sure to check the VMware website because we are currently in process of revalidating the building blocks with new guidance and pushing user count even higher.

VMware Exchange building blocks allow you to be very flexible in your deployment. For example, if you have 10,000 users, you could use two 4,000-user building blocks and one 2,000-user building blocks or five 2,000-user building blocks.

You can deploy using any combination of building blocks and still achieve predictable performance for your Exchange environment. No more over-provisioning your physical server that runs at 30 percent or less. By stacking virtual machines, you can get the full return on investment (ROI) by pushing that utilization to 50 percent or more.

Virtualizing Exchange Server 2010 Proof of Concept

With the success of virtualization of Windows, Active Directory, and IIS servers, Source Discovery has given much consideration to virtualizing its Exchange messaging services. The messaging organization is currently running on Exchange Server 2003 server clusters, and Source Discovery would like to transition its messaging organization to an Exchange 2010 messaging environment.

The company wants to virtualize the entire Exchange 2010 messaging environment. Currently, it's providing messaging and collaboration services for 5,000 users. Email archiving is also in use in the current messaging environment. Source Discovery has also stated that the transition of email archiving in the current environment will become a follow-on project to be implemented later.

The existing messaging environment consists of three Exchange Server 2003 failover clusters for high availability and to maximize their uptime within their organization. One failover cluster is actually dedicated to 200 executives and lawyers. In this environment, it is imperative that email remain operational at all times. The existing service-level agreement (SLA) is very close to five 9s or greater.

The other two clusters are dedicated to what they call tier-2 and tier-3 users. In this environment, the company has existing SLAs in place for these users. In the event of some business disruptive event, email operation must be restored in two hours or less with no more than one hour of mail data loss.

Email hygiene is handled via two load-balanced spam/firewall appliances. For the initial deployment, the messaging hygiene appliances will remain in place.

Source Discovery utilizes the most current released version of Microsoft Office 2007 with Service Pack 2 installed for all messaging clients. They also support Outlook Web Access or Outlook Web App, Outlook Mobile Access and Entourage, and Outlook Anywhere in the messaging environment.

Source Discovery has decided that it will replace the spam/firewall appliances with a virtualized Forefront threat management gateway solution after the full Exchange 2010 deployment has been completed.

With all the enhancements around the Microsoft Exchange 2010 release, Source Discovery has been given the budget to transition to Exchange 2010 along with the mandate to totally virtualize the deployment.

Working with Source Discovery Exchange administrators and systems administrators, you have decided to deploy Exchange 2010 with tier-1 and tier-2 users. Because the executives have an existing five 9s or greater availability requirement or SLA and the tier-2 users have a much less stringent requirement, you will deploy Exchange 2010 in two environments.

In the tier-1 environment, you will deploy using DAGs, and for the tier-2 users you will deploy them using stand-alone Exchange 2010 servers. Backups will be performed in the stand-alone server environment.

Although Exchange 2010 supports deployment on just a bunch of disks (JBOD), direct-attached storage (DAS), and storage area network (SAN) storage, you have decided to leverage the existing Source Discovery SAN storage for the new Exchange 2010 Server deployment.

In this proof of concept, we will cover the transition of the tier-1 users to the Exchange 2010 Server DAG environment. You will accomplish the following tasks:

1. Create the virtual machines for the server roles using templates for easier provisioning.
2. Install the Exchange prerequisites to prepare the virtual machines for the Exchange 2010 installation.

3. Install Exchange on all the server roles to exclude the Edge and Unified Messaging server roles.
4. Test the proposed solution using the Exchange LoadGen 2010 Beta tool on a isolated copy of the sourdisc.com Active Directory forest.

You will be integrating the Exchange 2010 server roles into the existing Exchange Server 2003 messaging environment. This will allow you to move mailboxes from the existing Exchange Server 2003 messaging environment to the newly created Exchange 2010 messaging environment. This process is called transitioning.

From here you will move into building the new Exchange 2010 Server environment.

Building the Exchange 2010 Server Environment

For the proof of concept, you will be building the Exchange 2010 DAGs environment for the tier-1 users. Because the DAG environment requires at least three Mailbox servers, Source Discovery has determined that it will reevaluate a portion of the more business-critical tier-2 users in the organization and upgrade these users to the tier-1 environment as needed.

This consideration will allow them to reduce the operational management overhead in the Exchange Server 2010 environment. By placing more users onto the tier-1 environment, the company can reduce the number of virtual machines to deploy in the stand-alone environment and maximize the storage in the Exchange Server 2010 DAG environment.

For the proof of concept, all Exchange Server 2010 machines will be joined to the isolated test sourdisc.com Active Directory forest to test the proposed Exchange Server 2010 solution using the Exchange LoadGen 2010 Beta tool. This environment was created beforehand from the production Active Directory forest and is available for testing.

After all tests have been completed, all Exchange 2010 databases will be deleted, and the Exchange Server 2010 virtual machines are placed in workgroup mode, the isolated Active Directory forest shut down. At this point, the Exchange Server 2010 virtual machines are rejoined to the live sourdisc.com forest. Both the production and isolated forest will be prepared for Exchange 2010 deployment.

First, you will check the existing Exchange Server 2003 environment to ensure that it meets all the prerequisites and checks before you introduce Exchange 2010 into the Exchange Server 2003 organization. This will involve the following:

1. Checking for or performing a current backup of the Exchange Server 2003 mailbox data in the organization
2. Ensuring that you are running Exchange Server 2003 on a supported version of the operating system and that all required service packs and updates are applied for the Exchange servers in the 2003 organization
3. Checking that the latest Exchange service packs and updates are applied on all Exchange servers in the 2003 organization

4. Verifying that the Exchange Server 2003 organization is in Exchange Server 2003 Native mode
5. Extending the Active Directory schema to receive the Exchange 2010 schema updates

Using Windows Server templates, you will build the various Exchange Server 2010 server roles' virtual machines to include installation into the Exchange Server 2003 messaging environment in the following order:

1. Exchange 2010 Client Access server
2. Exchange 2010 Hub Transport server
3. Exchange 2010 Mailbox server

After provisioning the virtual machines, you will then install the Exchange 2010 Server prerequisites, and finally you will install Exchange 2010 Server onto each virtual machine.

You will then create the DAG environment and provision the database copies. There will be three copies of the mailbox databases in this environment.

From there you will cover some of the commands to check the state of database replication to ensure that the Source Discovery deployment is operating as expected.

You will not be installing the following roles in this proof of concept:

- The Unified Messaging server role (UM), because it is not currently supported for virtualization by Microsoft
- The Edge transport role (this will be installed in a follow-on project to replace the messaging hygiene appliances)

After verifying the current environment, Source Discovery is ready to proceed with deploying the virtual machines from templates. You will first deploy the Client Access server (CAS) role using a created template.

Preparing the Active Directory Environment for Exchange Server 2010

To introduce the Exchange Server 2010 roles into the existing Exchange Server 2003 messaging environment, you must first prepare the existing Active Directory environment to receive the new schema classes and attributes required for proper Exchange Server 2010 deployment and operation. After ensuring that Active Directory is in the proper operation mode for Exchange Server 2010 deployment, you can begin to update and extend Active Directory.

In Chapter 4, you learned how to virtualize Active Directory, and you deployed in Windows 2003 forest and domain operation mode. You will confirm that you are in Windows Server 2003 forest and domain operation mode. Figure 6.7 confirms Active Directory operational mode. In Windows Server 2008 R2, you now can use the Active Directory Administrative Center (ADAC) to determine the forest and domain operating modes. To open the ADAC, from the Start button on the desktop, click Administrative Tools and select Active Directory Administrative Center from the menu.

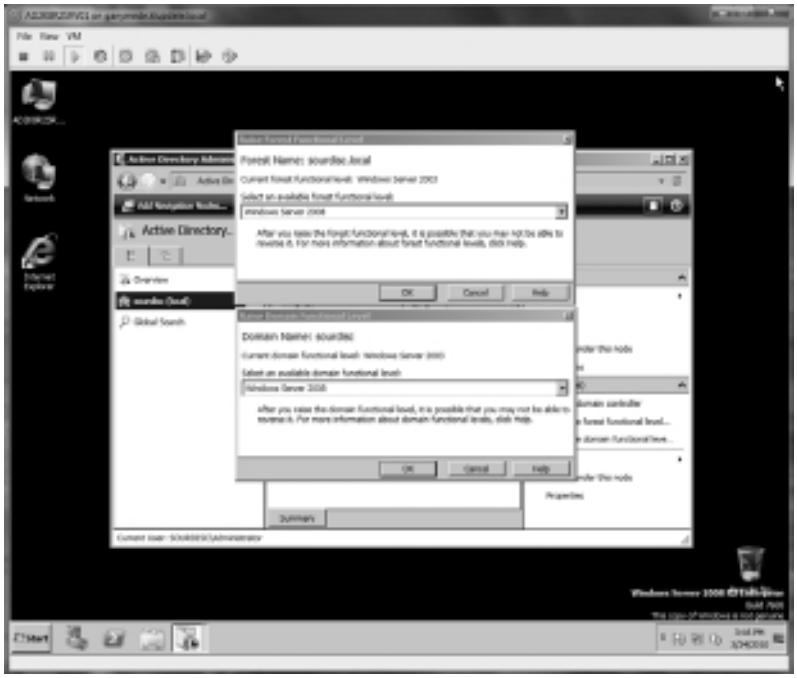

FIGURE 6.7
Confirming Active Directory forest and domain operational mode

Once you have identified the Active Directory forest and domain operational mode, you must identify the Exchange Server 2003 operational mode. Open the Exchange System Manager Start menu and Exchange Server 2003 menu, and right-click the Exchange organization.

In the operation mode field, ensure that Exchange Server 2003 is operating in the Native (not pre–Exchange 2000 servers) mode. When Exchange is operating in this mode, no more previous-version Exchange servers can be installed into the organization. Figure 6.8 shows the Exchange Server 2003 operational mode.

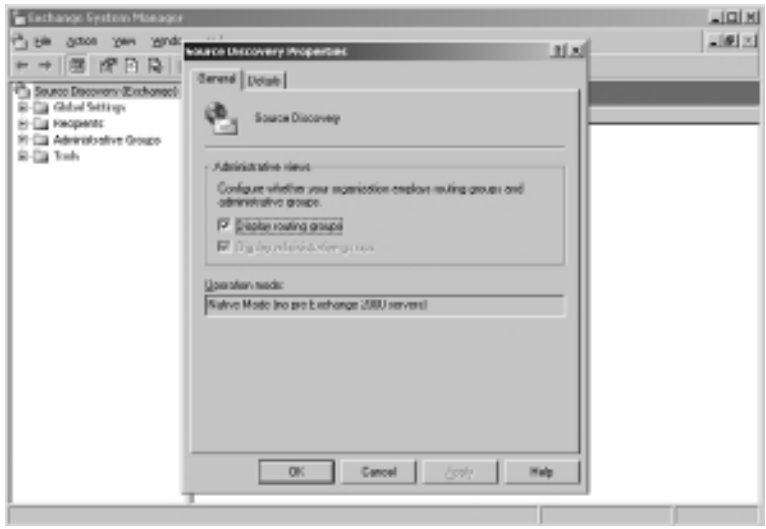

FIGURE 6.8
Confirming Exchange operating mode

Now that you have verified that Active Directory is in Windows Server 2003 operation mode, you can prepare Active Directory to receive Exchange 2010.

If you are installing Exchange Server 2010 from a privileged user account, one that is a member of the Enterprise Administrators, Domain Administrators, and Schema Administrators groups, the update to the Active Directory will be performed without manual intervention.

In the proof of concept, we will illustrate how to prepare Active Directory for Exchange Server 2010. The tasks for preparing Active Directory include the following:

1. Preparing the schema
2. Preparing Active Directory
3. Preparing the domain

These tasks are described in detail in the following sections.

Preparing the Active Directory Schema

Before installing Exchange Server 2010 into your existing Exchange Server 2003 messaging environment, the Active Directory forest that will host Exchange Server 2010 will have to have its schema extended.

Extending the schema will allow Active Directory to facilitate all the features and functionalities included in the new version of Exchange. Extending the schema introduces the new classes and attributes that would be required for proper Exchange Server 2010 operation.

You open a command window and change to the Exchange Server 2010 media. Using the following command, you can prepare the Active Directory schema:

```
Setup /PrepareSchema
```

Figure 6.9 illustrates the steps in preparing the Active Directory schema.

Administrator: Command Prompt
R:\>Setup.com /PrepareSchema
Welcome to Microsoft Exchange Server 2010 Unattended Setup
By continuing the installation process, you agree to the license terms of Microsoft Exchange Server 2010. If you don't accept these license terms, please cancel the installation. To review these license terms, please go to http://go.microsoft.com/fwlink/?LinkId=150127&clcid=0x409
Press any key to cancel setup.....
No key presses were detected. Setup will continue.
Preparing Exchange Setup
Copying Setup Files ..... COMPLETED
No server roles will be installed
Performing Microsoft Exchange Server Prerequisite Check
Organization Checks ..... COMPLETED
Configuring Microsoft Exchange Server
Extending Active Directory schema
Progress: ..... COMPLETED
The Microsoft Exchange Server setup operation completed successfully.
R:\>

Figure 6.9
Active Directory schema preparation

As you can see from the illustration, the command completed successfully. You should not proceed any further if this step fails.

Preparing the Active Directory

Moving on, you will now prepare the Configuration partition in Active Directory to receive Exchange Server 2010. As we mentioned earlier in this section, you must be a member of the Enterprise Administrators group because the Domain Administrator can only read and not write to the Configuration partition on the Active Directory.

If you were installing a "green" or a new Exchange Server 2010 messaging organization, you would be required to use the `/Organization:` parameter on the `Setup` command.

You would then receive a warning message stating that previous Exchange Servers would not be able to participate in the new Exchange Server 2010 organization.

NOTE When installing a new Exchange Server 2010 messaging organization and using the `Setup /PrepareAD /Organization:` parameter, the following message will be displayed: "After this operation, you will not be able to install any Exchange 2007 server roles."

To prepare the Active Directory, you enter the following command from the command line:

```
Setup /PrepareAD
```

This will update the Active Directory Configuration partition with the necessary information to receive Exchange Server 2010 into the existing Exchange Server 2003 messaging organization. Figure 6.10 shows the `Setup /PrepareAD` command completed successfully. The Active Directory has been successfully updated.

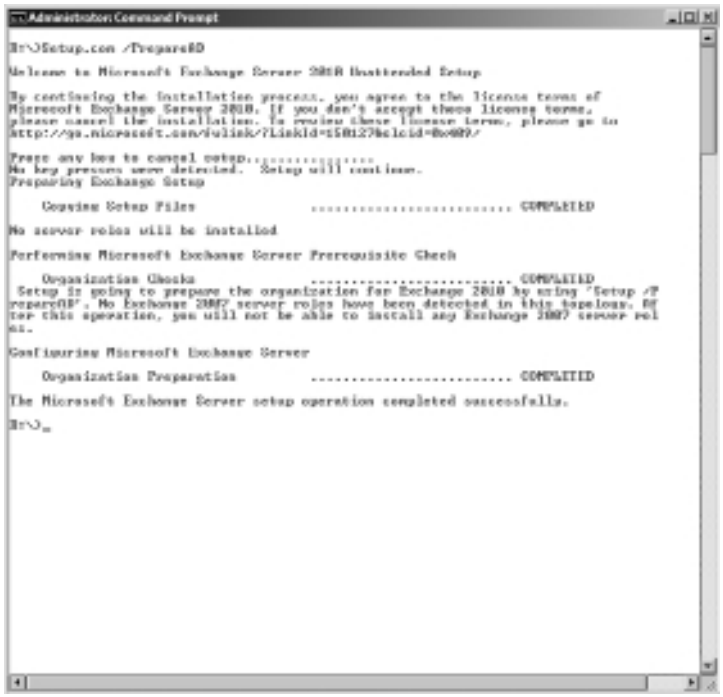
Administrator: Command Prompt
D:\>Setup.com /PrepareAD
Welcome to Microsoft Exchange Server 2010 Unattended Setup
By continuing the installation process, you agree to the license terms of
Microsoft Exchange Server 2010. If you don't accept these license terms,
please cancel the installation. To review these license terms, please go to
http://go.microsoft.com/fwlink/?LinkId=150127&clcid=0x409.
Press any key to cancel setup.....
No key presses were detected. Setup will continue.
Preparing Exchange Setup
Copying Setup Files ..... COMPLETED
No server roles will be installed
Performing Microsoft Exchange Server Prerequisite Check
Organization Checks ..... COMPLETED
Setup is going to prepare the organization for Exchange 2010 by using 'Setup /P
repareAD'. No Exchange 2007 server roles have been detected in this topology. Af
ter this operation, you will not be able to install any Exchange 2007 server rol
es.
Configuring Microsoft Exchange Server
Organization Preparation ..... COMPLETED
The Microsoft Exchange Server setup operation completed successfully.
D:\>

FIGURE 6.10
Updating Active Directory

The final step in preparing the Active Directory to receive Exchange Server 2010 into the current messaging environment is the `Setup /PrepareDomain` command. This command must be issued for any Active Directory domain that will host Exchange Server 2010. You can issue this command with the `/PrepareAll Domains` parameter to prepare all the domains in the Active Directory forest.

When you issue this command, the domain is prepared to receive Exchange Server 2010. We will issue the following command from the command line:

```
Setup /PrepareDomain
```

Figure 6.11 displays the output of the prepare domain command.

FIGURE 6.11
Preparing the Active Directory domain

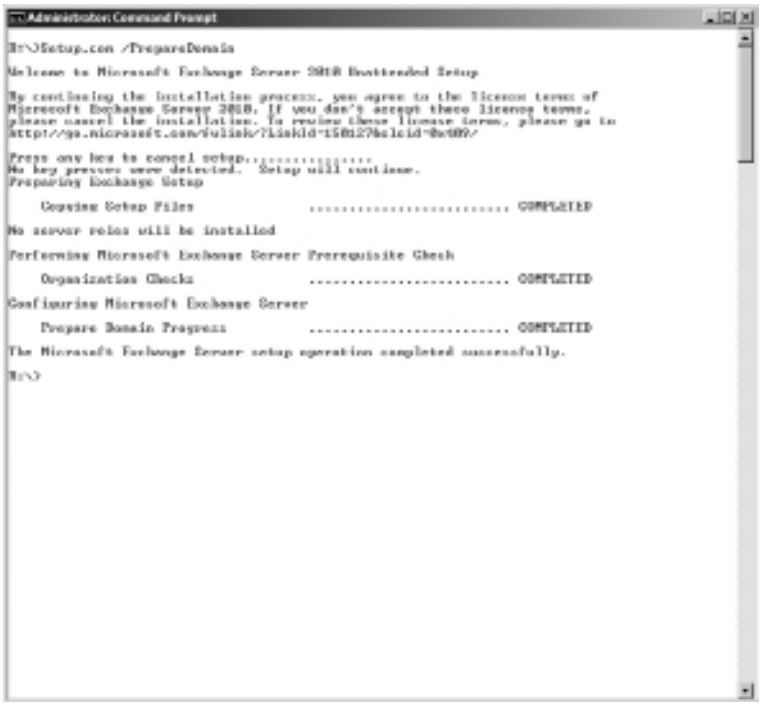
Administrator: Command Prompt
E:\>Setup.com /PrepareDomain
Welcome to Microsoft Exchange Server 2010 Unattended Setup
By continuing the installation process, you agree to the license terms of
Microsoft Exchange Server 2010. If you don't accept these license terms,
please cancel the installation. To review these license terms, please go to
http://go.microsoft.com/fwlink/?LinkId=150127&clcid=0x409/
Press any key to cancel setup.....
No key presses were detected. Setup will continue.
Preparing Exchange Setup
Copying Setup Files ..... COMPLETED
No server roles will be installed
Performing Microsoft Exchange Server Prerequisite Check
Organization Checks ..... COMPLETED
Configuring Microsoft Exchange Server
Prepare Domain Progress ..... COMPLETED
The Microsoft Exchange Server setup operation completed successfully.
E:\>

As you can see from Figure 6.11, the command completed successfully. In the process of this preparation, a number of security groups are created in a root-level container in the domain called Microsoft Exchange Security Groups. These are the built-in roles or security groups that can be used to delegate Exchange Server 2010 administration roles. Figure 6.12 displays the Microsoft Exchange Security Groups container created after you prepared the Source Discovery organization.

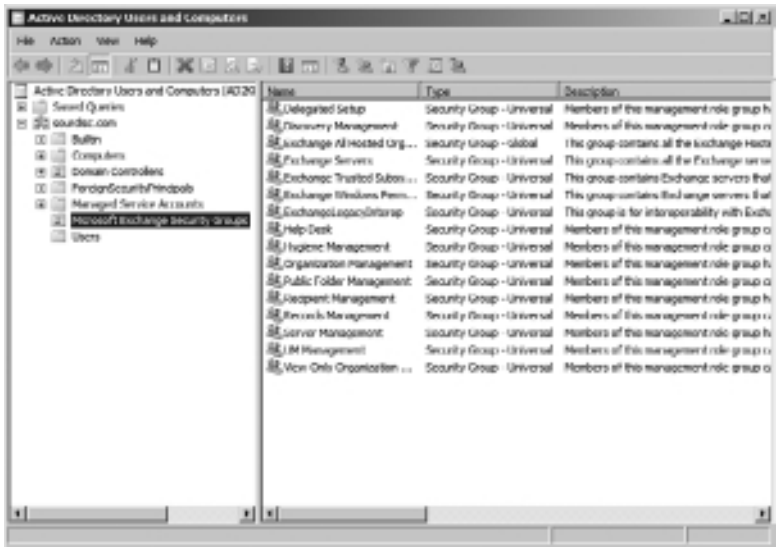

FIGURE 6.12
Microsoft Exchange Security Groups

Now that you have prepared the active domain for Exchange Server 2010, it is time to build the CAS server virtual machine, install the Exchange Server 2010 server role prerequisites, and install Exchange Server 2010 onto the virtual machine.

Building the Exchange Client Access Server Role Virtual Machine

Using a previously created virtual machine template, we will provision and deploy the first CAS role into the Exchange Server 2003 messaging environment.

Using virtual machine templates will radically reduce the time required to provision the Exchange Server 2010 roles in the existing Exchange Server 2003 messaging environment.

Using vCenter, select Home ➢ Inventory ➢ VMs And Templates, select the virtual machine template, and click Deploy Virtual Machine From This Template. This will start the Deploy From Template Wizard. In the wizard, you will enter the following information:

- The name of the virtual machine
- The VMFS data store on which to deploy the virtual machine
- The disk format (whether to deploy using thick or thin format)
- Whether you want to customize the virtual machine
- Any additional information required to customize the VM if guest customization is to be used

Figure 6.13 shows the Deploy From Template Wizard. As you move through the screens within the wizard, you will enter or choose the options that are most appropriate for your deployment of the template.

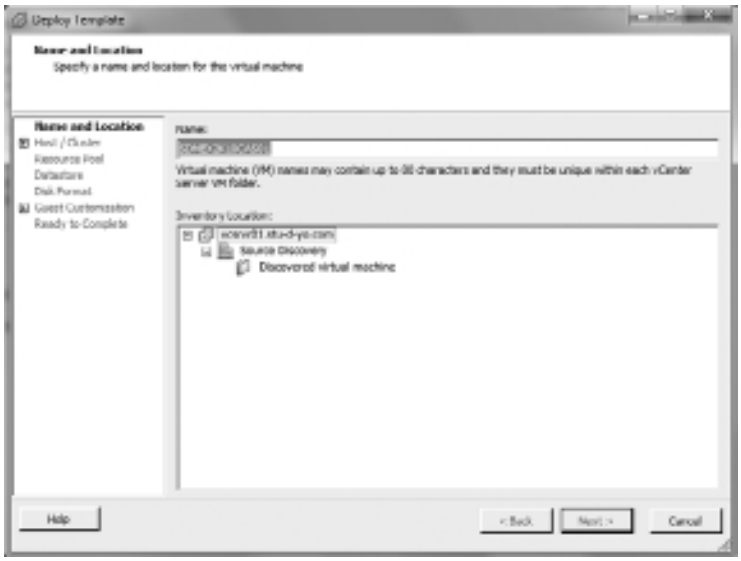

FIGURE 6.13
Deploy From Template Wizard screen

After you have entered all the necessary information, choose to use guest customization. Guest customization allows you to specify the virtual machine properties such as the following:

- Registration information
- Operating system NetBIOS name
- Windows operating system product key and the server licensing mode
- Password for the Administrator account
- Time zone information
- Any commands that you would want to run once upon startup of VM
- Network settings
- Selection of workgroup or domain mode
- Choice of whether to generate a new security identifier (SID)

Once the selections are completed, you are presented with the Ready To Complete screen. This is the last chance to correct any selection before actually deploying the new virtual machine from the chosen template.

From this point, you click the Finish button to deploy the virtual machine from the selected virtual machine template. Figure 6.14 displays the guest customization Ready To Complete screen in which you can view your selections and modify them if needed before you deploy the virtual machine from the virtual machine template.

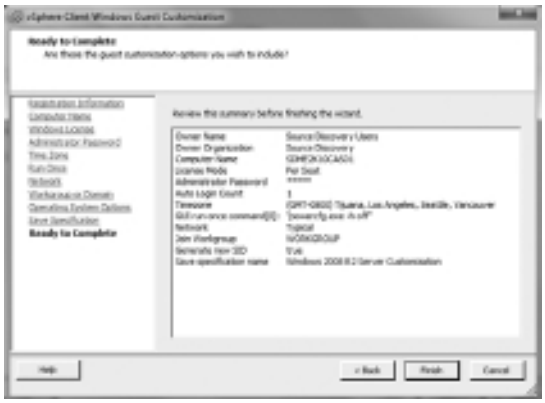

FIGURE 6.14 Confirming selections before deployment

Once the virtual machine has been deployed, you can continue the process of installing the Exchange Server 2010 prerequisites on the virtual machines. We have not included Exchange 2010 prerequisites in the template so that we can illustrate the process of installing the required features and roles for Exchange Server 2010.

Once the virtual machine has booted and is joined to the domain, you can now continue with the Exchange Server 2010 prerequisites. When installing the Exchange 2010 prerequisites for the CAS role, you have the option of manually installing the roles and features from the command line. Table 6.19 lists the prerequisites for the Exchange 2010 CAS role.

TABLE 6.19: Exchange 2010 CAS Server Role Prerequisites

.NET Framework 3.5	PowerShell v2.0
Windows Remote Management	Hot Fix 951725 (applies to Windows 2008)
Web Server	Web Server Metabase
Web Server Legacy Mgmt Console	Web Server Basic Authentication
Web Server Windows Authentication	Web-Net-Extensions
Web-Digest-Authentication	Web-Dynamic-Compression
Web-ISAPI-Extensions	RPC-Over-HTTP-Proxy (Outlook Anywhere)
Remote Server Administrator Tools (required to facilitate schema changes)	Net-HTTP-Activation

WARNING The proper versions of .NET, PowerShell, and Windows Remote Management are built in to Windows Server 2008 R2. If you are installing the prerequisites on Windows Server 2008, ensure you download and install the proper versions of these programs available from the Microsoft website.

If you decide to manually install these, ensure that you understand all the requirements for the CAS role. You can install all the features and roles by typing all the requirements from the command line. Type the following:

```
ServerManagerCmd -i Web-Server Web-ISAPI-Ext Web-Metabase Web-Lgcy-Mgmt-Console -
Web-Basic-Auth Web-Digest-Auth Web-Windows-Auth Web-Dyn-Compression Web-Net-Ext -
RPC-over-HTTP-proxy Net-HTTP-Activation
```

Figure 6.15 shows the command line for installing the Exchange 2010 server role prerequisites.

Administrator: Command Prompt
C:\>ServerManagerCmd -i Web-Server Web-ISAPI-Ext Web-Metabase Web-Lgcy-Mgmt-Cons
ole Web-Basic-Auth Web-Digest-Auth Web-Windows-Auth Web-Dyn-Compression Web-Net-
Ext RPC-over-HTTP-proxy Net-HTTP-Activation_

FIGURE 6.15
Installing the server role prerequisites from the command line

For those of you who prefer to avoid all of the typing, Microsoft has given you a more efficient method of installing all the Exchange 2010 prerequisites.

Microsoft has created XML files that you can use with command line of the Server Manager to get the Exchange 2010 server role prerequisites installed in a more efficient manner. These XML files are located in the `Scripts` directory on the Exchange 2010 media. These files cover the typical installation; we will talk more about a typical installation later when we talk about the custom installation of the Exchange 2010 server roles. Figure 6.16 illustrates the Microsoft-provided XML files located in the `Scripts` directory on the Exchange 2010 media.

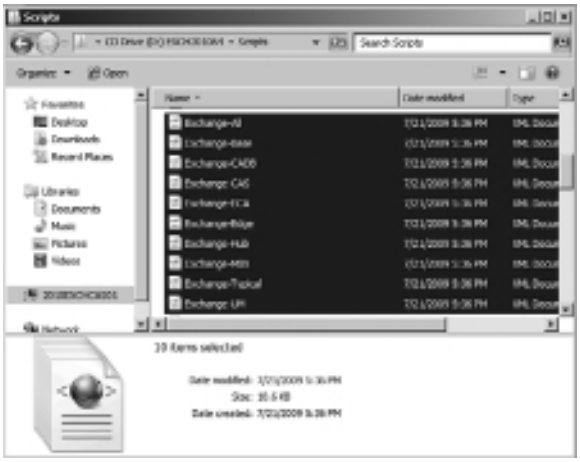

FIGURE 6.16
Exchange 2010 `Scripts` directory

As you can see, there is an `.xml` file for all the server roles and also a separate file for each Exchange 2010 server role.

You will use the Server Manager command line to install the CAS server role prerequisites onto the CAS virtual machine. At the command line, type the following command:

```
C:\ ServerManagerCmd -InputPath or -ip Media:\Scripts\Exchange-CAS.xml
```

This will initiate the installation of the CAS server role prerequisites onto the CAS virtual machine. Figure 6.17 shows the installation of the Exchange Server 2010 CAS server role using `Exchange-CAS.xml` from the command line via the Server Manager command-line mode.

Administrator: Command Prompt
C:\>ServerManagerCmd.exe -InputPath E:\Scripts\Exchange-CAS.xml
Servermanagercmd.exe is deprecated, and is not guaranteed to be supported in fut
ure releases of Windows. We recommend that you use the Windows PowerShell cmdlet
s that are available for Server Manager.
Start Installation...
[Installation] Succeeded: [.NET Framework 3.5.1 Features] .NET Framewor
k 3.5.1. You must restart this server to finish the installation process.
[Installation] Succeeded: [Windows Process Activation Service] Configur
ation APIs. You must restart this server to finish the installation process.
[Installation] Succeeded: [Windows Process Activation Service] .NET Env
ironment. You must restart this server to finish the installation process.
[Installation] Succeeded: [Windows Process Activation Service] Process
Model. You must restart this server to finish the installation process.
[Installation] Succeeded: [Remote Server Administration Tools] Failover
Clustering Tools. You must restart this server to finish the installation proce
ss.
[Installation] Succeeded: [Remote Server Administration Tools] Web Serv
er (IIS) Tools. You must restart this server to finish the installation process.
[Installation] Succeeded: [Web Server (IIS)] IIS Management Console. Yo
u must restart this server to finish the installation process.
[Installation] Succeeded: [Web Server (IIS)] Static Content. You must r
estart this server to finish the installation process.
[Installation] Succeeded: [Web Server (IIS)] Tracing. You must restart
this server to finish the installation process.

FIGURE 6.17
Installing the Exchange 2010 prerequisites using Exchange 2010 scripts

Let's touch on what a typical installation is comprised of:

- Client Access server role
- Hub Transport role
- Mailbox role
- Exchange 2010 Server Management Tools

In a typical installation of the Exchange Server 2010 roles, you would install the CAS, HUB, MBX, and management tools onto a single server. This is an excellent way for smaller organizations to save money on server hardware and licensing for their deployments of Exchange Server 2010 by reducing the hardware and using the Exchange Server 2010 standard licensing.

Just because a smaller organization might deploy Exchange Server 2010 on a single server does not mean that they are vulnerable to the single-server failure scenario. Exchange Server 2010 now provides the means for deploying an end-to-end highly available solution using just two virtual machines on two separate ESX hosts.

You can combine the typical installation into a DAG scenario and create a highly available solution for a small to medium business using Exchange Server 2010 standard licensing.

In the custom installation, you could install any single or combination of the Exchange Server 2010 server roles, excluding the Edge server role because the Edge role must be installed on a separate virtual machine in the DMZ of your organization's network.

Figure 6.18 displays the Exchange Server 2010 installation options.

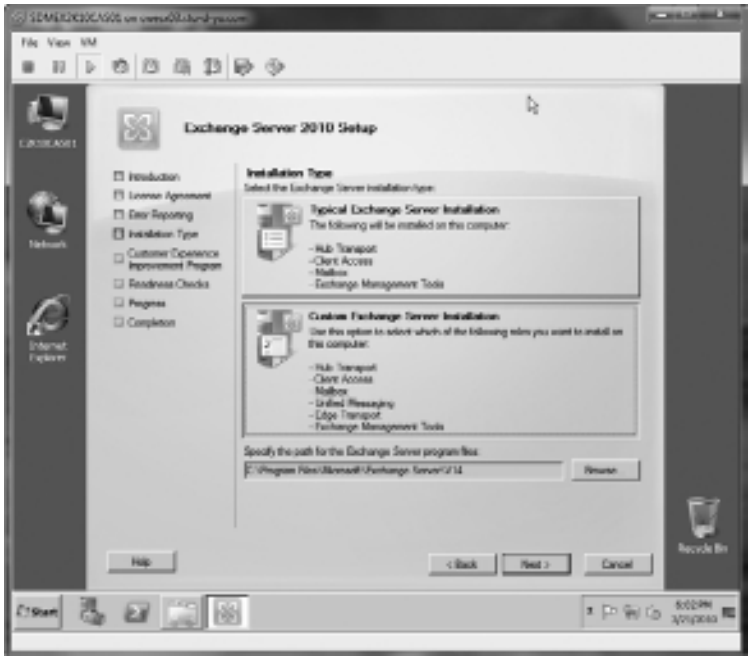

FIGURE 6.18
Exchange Server 2010 installation options

As you can see from Figure 6.18, under the custom installation option, the following choices are available:

- Client Access role
- Hub Transport role
- Mailbox role
- Unified Messaging role
- Edge Transport role
- Exchange 2010 Server Management Tools

In the Source Discovery deployment, we will be using the custom mode of installation because we will be installing the Exchange Server 2010 roles onto separate virtual machines.

Deploying the Exchange Server 2010 CAS Role

Now that we have confirmed the Active Directory operating mode, prepared Active Directory to receive Exchange Server 2010, built the CAS server role virtual machine, and installed the Exchange Server 2010 prerequisites, we can now install the Exchange Server 2010 CAS roles into the existing Exchange Server 2003 messaging environment.

TIP Before starting the Exchange Server 2010 install on the CAS server role, be sure to set the Net.Tcp Port Sharing Service option to Automatic in the Windows Services console.

Installing the Exchange Server 2010 Client Access server role is as easy as popping the DVD into the drive running the Setup program and choosing the custom installation.

You connect the virtual machine to Exchange Server 2010 media, wait for the new media in drive notification, and choose to run the Setup program. If the autorun program does not start automatically, you need to go to the root of the Exchange Server 2010 media and double-click the Setup program to start it. Once the Setup program has started, you will receive the splash screen. Figure 6.19 shows the Exchange Server 2010 splash screen.

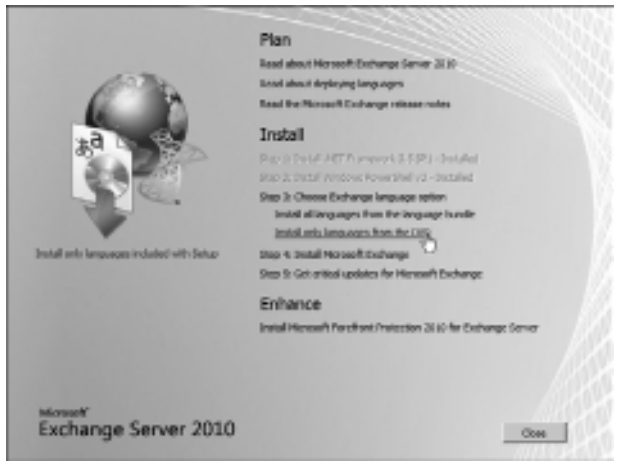

Figure 6.19 Starting the Exchange Server 2010 Setup program

Notice how the PowerShell and the .NET 3.5 steps are grayed out and we can move only to step 3 of selecting the languages to install. The Exchange Server 2010 Setup program actually uses this method of steps to prevent you from skipping any necessary steps in the installation process.

If we select the Install Exchange option without choosing the languages step, the setup program will present you with an error message that states that we must choose the languages to install. Here we have chosen to install only the languages present on the Exchange Server 2010 installation media.

After installing the languages from the DVD, you will see in the next illustration that the Install Exchange option has now become step 4. From here, we can choose to install Exchange Server 2010 on the CAS virtual machine. In Figure 6.20, notice how the Install Exchange option has become a step.

FIGURE 6.20
Installing Exchange Server 2010

From here we choose Install and the installation proceeds to the introduction, end-user licensing agreement (EULA), error reporting, and the installation type screens, which is where we will choose the typical or custom installation option.

For our proof of concept, choose the custom installation in order to install only the Client Access server role and the Exchange Server 2010 management tools on the CAS server virtual machine. Refer to Figure 6.18 to see the installation options for Exchange Server 2010.

In Figure 6.21 you will notice that we chose the CAS server role for installation. Also notice that the Exchange Management Tools were also selected and grayed out.

FIGURE 6.21
CAS server role installation

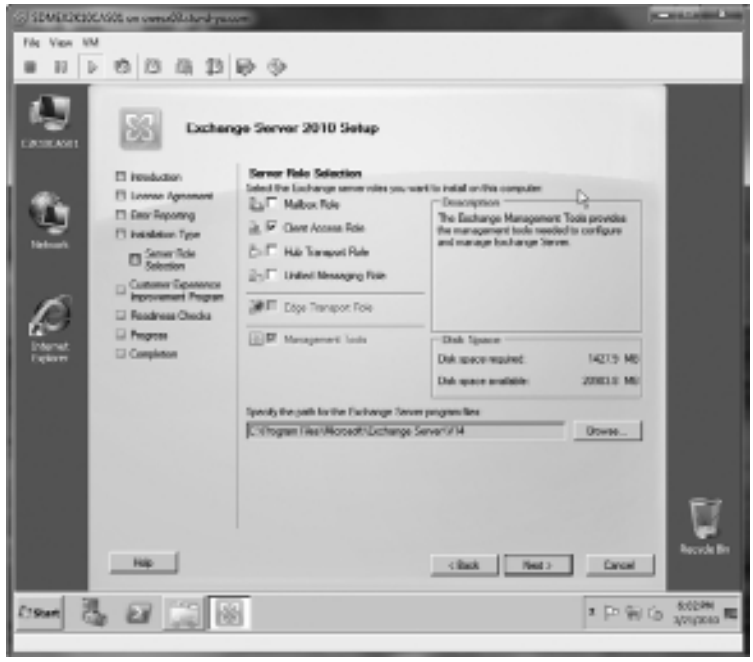

TIP If this were a clean install of Exchange 2010 Server into a new messaging environment, you would be prompted to configure an external domain, meaning that the Client Access servers would be Internet-facing.

During the installation process, the Setup program checks that all prerequisites are present and installed. Other readiness checks such as picking a language, copying the Exchange Server 2010 files, installing the Exchange Server 2010 Management Tools, and installing the CAS server role are all done during this process.

Figure 6.22 shows the successful completion of the CAS server role installation. Any installation errors that occur must be corrected before continuing with any other actions.

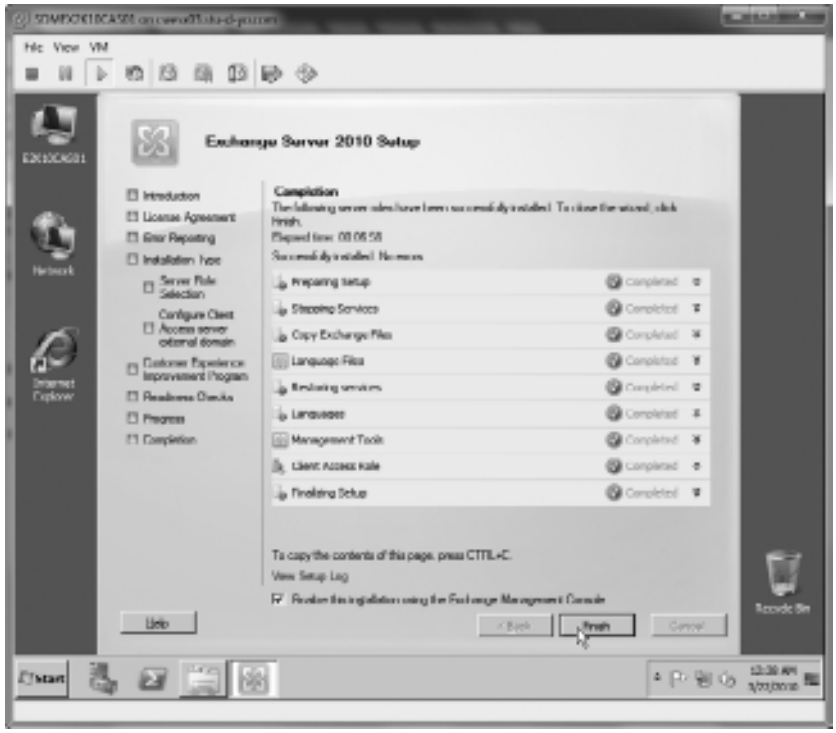

FIGURE 6.22 Successful CAS server role completion

After installing the CAS server role, you should consider doing some of these post-installation tasks:

1. Enter the Exchange 2010 Server product key.
2. Configure secure access to the CAS server.
3. Configure ActiveSync and OWA access.

Typically, these tasks will be performed by the Source Discovery Exchange and system administrators in accordance with the company's operational and security policies.

TIP For more information on the post-installation completion tasks, see the Microsoft TechNet website:

```
http://technet.microsoft.com/en-us/library/bb125262.aspx?=14.0.639.21&t=exchgf1
```

You will find the latest information around all the tasks for finalizing your Exchange Server 2010 deployment. It should be stated that just because you are virtualizing Exchange Server 2010 does not mean you can omit any of the required installation processes and procedures.

With the Exchange Server 2010 CAS server role installed into the Source Discovery Exchange Server 2003 messaging environment, we will move on to deploying the Hub Transport server role.

Deploying the Exchange Server 2010 Hub Transport Server Role

Using the steps from the Exchange Server 2010 CAS server role deployment, we will deploy the Hub Transport into the Source Discovery Exchange Server 2003 messaging environment.

With the Hub Transport virtual machine deployed from the existing template, we will install the Exchange Server 2010 Hub Transport role prerequisites. Table 6.20 lists the prerequisites for the Exchange Server 2010 Hub Transport server role.

Table 6.20: Exchange 2010 Hub Transport Server Role Prerequisites

.NET Framework 3.5	PowerShell v2.0
Windows Remote Management	Hot Fix 951725 (applies to Windows 2008)
Web Server	Web Server Metabase
Web Server Legacy Mgmt Console	Web Server Basic Authentication
Web Server Windows Authentication	Web-Net-Extensions
Remote Server Administrator Tools (required to facilitate schema changes)	
Microsoft Office 2007 Converter Filter Pack	

Using the Exchange Server 2010 scripts that come on the installation media, you can install all the Exchange Server 2010 Hub Transport role prerequisites by using the command-line version Server Manager in a command window to issue the following command:

```
ServerManagerCmd -InputPath or -ip Media:\Scripts\Exchange-HUB.xml
```

Running this command will install all the Exchange Server 2010 Hub Transport prerequisites on the Hub Transport virtual machine. Figure 6.23 shows the output of the Server Manager command.

With the Hub Transport prerequisites installed, you must now install the Microsoft Office 2007 Filter Pack application on the Hub Transport virtual machine. You can download this application from the Microsoft website.

Administrator: Command Prompt
C:\>ServerManagerCmd -inputPath R:\Scripts\Exchange-Hub.xml
Servermanagercmd.exe is deprecated, and is not guaranteed to be supported in fut
ure releases of Windows. We recommend that you use the Windows PowerShell cmdlet
s that are available for Server Manager.
Start Installation...
Warning: [Installation] Succeeded: [.NET Framework 3.5.1 Features] .NET Framewor
k 3.5.1. You must restart this server to finish the installation process.
Warning: [Installation] Succeeded: [Windows Process Activation Service] Configur
ation APIs. You must restart this server to finish the installation process.
Warning: [Installation] Succeeded: [Windows Process Activation Service] .NET Env
ironment. You must restart this server to finish the installation process.
Warning: [Installation] Succeeded: [Windows Process Activation Service] Process
Model. You must restart this server to finish the installation process.
Warning: [Installation] Succeeded: [Remote Server Administration Tools] Web Serv
er (IIS) Tools. You must restart this server to finish the installation process.
Warning: [Installation] Succeeded: [Web Server (IIS)] IIS Management Console. Yo
u must restart this server to finish the installation process.
Warning: [Installation] Succeeded: [Web Server (IIS)] Default Document. You must
restart this server to finish the installation process.
Warning: [Installation] Succeeded: [Web Server (IIS)] Static Content. You must r
estart this server to finish the installation process.
Warning: [Installation] Succeeded: [Web Server (IIS)] Static Content Compression
. You must restart this server to finish the installation process.
Warning: [Installation] Succeeded: [Web Server (IIS)] IIS 6 Metabase Compatibili
ty. You must restart this server to finish the installation process.
Warning: [Installation] Succeeded: [Web Server (IIS)] Request Filtering. You mus
t restart this server to finish the installation process.

FIGURE 6.23
Hub Transport prerequisite installation

After installing the Filter Pack application, you can now proceed to install the Exchange Server 2010 Hub Transport server role onto the Hub Transport virtual machine. Start the Setup program to begin installing the Exchange Server 2010 Hub Transport role.

Figure 6.24 will show a process similar to that which you followed while installing the Exchange Server 2010 CAS server role.

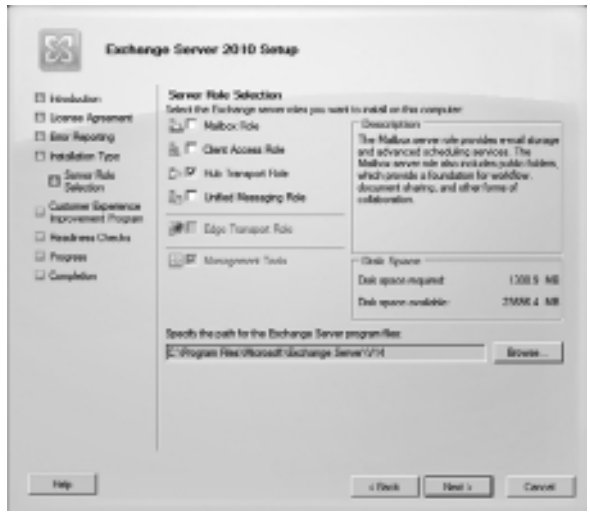

FIGURE 6.24
Installing the Hub Transport server role

In the figure, the Hub Transport server role is selected along with the Exchange Server 2010 management tools. When you proceed to the next screen, we will be asked to choose an Exchange server that can provide routing functions for the Exchange Server 2003 environment.

In Figure 6.25 we choose the server that will perform the routing functions for the Exchange Server 2003 messaging environment.

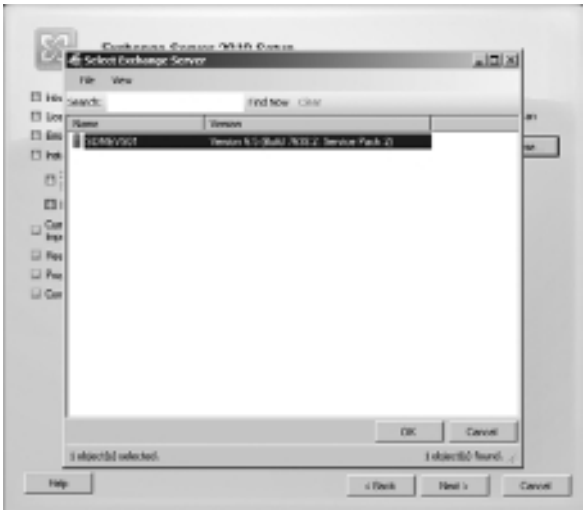

FIGURE 6.25
Selecting the Exchange Server 2003 routing server

After performing the readiness checks, copying the Exchange files, installing the Exchange Server management tools, and installing the Hub Transport server role on to the Hub Transport virtual machine, the installation program completes and prompts you to get any new updates from the Internet.

You should always choose step 5 to download and install the new updates from the Internet because this will ensure that you are running the latest version of the Exchange 2010 program with the latest fixes and updates that have become available after the initial release of the Exchange Server 2010 message application.

In Figure 6.26, we see the successful completion of the Exchange installation of the Hub Transport server role.

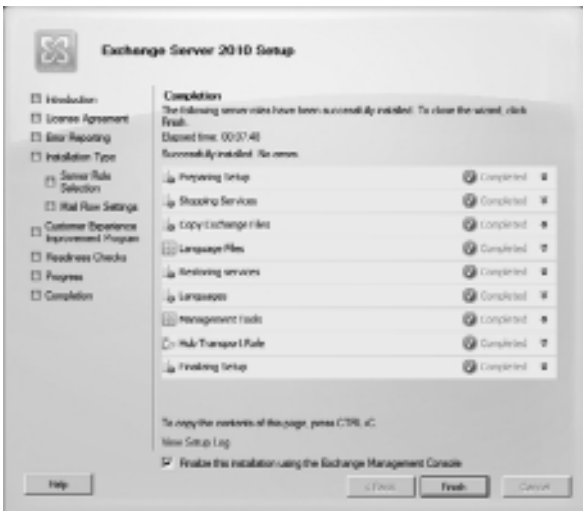

FIGURE 6.26
Completion of the Hub Transport server role

After installing the Hub Transport server role, we can route messages to and from the Exchange Server 2003 and Exchange Server 2010 messaging environments. From here we will move on to installation of the Exchange Server 2010 Mailbox server role into the Source Discovery Exchange Server 2003 messaging environment.

Deploying the Exchange Server 2010 Mailbox Server Role

When deploying the Exchange Server 2010 Mailbox server role into the Source Discovery Exchange Server 2003 messaging organization, we will have completed the basic installation of the Exchange Server 2010 messaging environment without the Edge Transport and Unified Messaging server roles.

In this example, we will install three Mailbox server roles into the current messaging organization to facilitate the creation of the DAG required for the Source Discovery executives and lawyers. Table 6.21 lists the required prerequisites for the Exchange Server 2010 Mailbox server role.

Table 6.21: Exchange 2010 Mailbox Server Role Prerequisites

.NET Framework 3.5	PowerShell v2.0
Windows Remote Management	Hot Fix 951725 (applies to Windows 2008)
Web Server	Web Server Metabase
Web Server Legacy Mgmt Console	Web Server Basic Authentication
Web Server Windows Authentication	Web-Net-Extensions
Remote Server Administrator Tools (required to facilitate schema changes)	Failover Clustering (only required for database availability)
Microsoft Office 2007 Converter Filter Pack	

For the Exchange Server 2010 Mailbox server role, you will notice the requirement for Windows Server 2008 failover clustering. It should be noted that this feature is required to be installed only when you will employ DAGs for the Exchange Server 2010 Mailbox server role. Because we are deploying a DAG for the VIPs in the Source Discovery organization, we will install this feature on all three of the Exchange Server 2010 Mailbox server roles.

Following the procedures for installing the Exchange Server 2010 Hub Transport server prerequisites, issue the following command from a command window:

```
ServerManagerCmd -InputPath or -ip Media:\Scripts\Exchange-MBX.xml
```

This will install all the prerequisites excluding the Microsoft Office 2007 Filter Pack application for the Mailbox server role virtual machines. Figure 6.27 shows the output of a successful installation of the Mailbox server role prerequisites on to the Mailbox server role virtual machine.

FIGURE 6.27
Mailbox server role prerequisites installation

After we install the prerequisites and reboot the virtual machine, we are ready to install the Microsoft Office 2007 Filter Pack application. After installing this last prerequisite, we can run the Exchange Server 2010 Setup program and install the Exchange Server 2010 Mailbox server role into the Source Discovery Exchange Server 2003 messaging organization.

Figure 6.28 shows the installation program running the readiness checks, and you encounter a warning about replicating the free/busy information to other free/busy servers within the organization after the installation of Exchange Server 2010 has been completed.

FIGURE 6.28
Installing the Exchange Server 2010 Mailbox server role

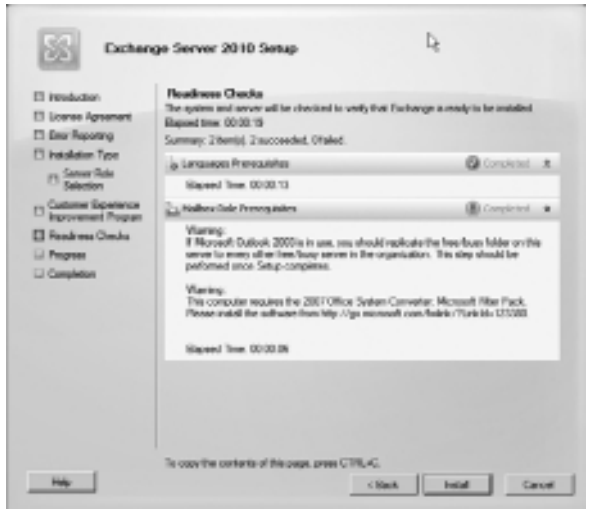

After the Setup program has performed all the steps in the MBX transport installation and there are no errors during the installation process, the installation completes successfully. At this point, we are ready to repeat the procedure for the two remaining Exchange Server 2010 Mailbox server role virtual machines. Figure 6.29 shows the successful completion of the Exchange Server 2010 setup program.

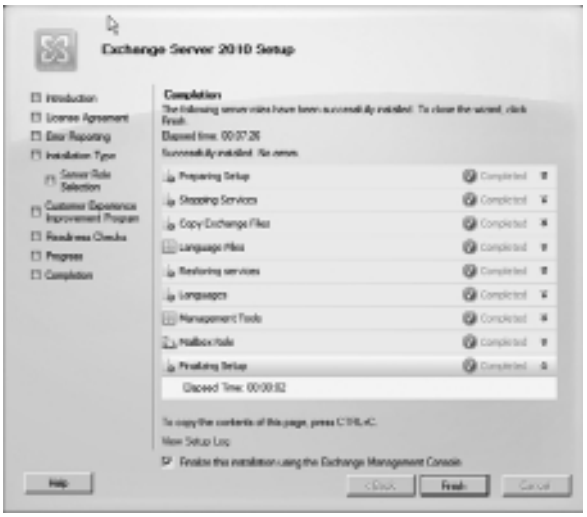

FIGURE 6.29
Exchange Server 2010 Mailbox server role installation completion

After building the Exchange 2010 infrastructure, we can create a separate Active Directory forest and join the virtual machines to this forest to perform a validation of the proposed Exchange Server 2010 solution using the Exchange Server 2010 Load Generator Beta tool.

NOTE We tested the storage subsystem to be used for Exchange Server 2010 with the JetStress tool. This is a Microsoft best practice and will ensure that your proposed storage solution will meet or exceed the requirements of the Exchange Server 2010 product.

TESTING THE PROPOSED EXCHANGE 2010 SOLUTION WITH LOADGEN 2010 BETA

With the Exchange Server 2010 virtual machines joined to the isolated sourdisc Active Directory forest, we can now proceed with testing the Exchange Server 2010 solution using the Exchange LoadGen 2010 Beta tool.

The Exchange Server LoadGen 2010 client server is installed after the Windows Server 2008 R2 operating system and the Web server role is installed. The server is joined to the isolated sourdisc.com Active Directory forest.

After a reboot of the server is performed, we then open the Exchange Server LoadGen 2010 Beta tool. In Figure 6.30 you see the LoadGen 2010 Beta test tool opening screen.

From here we can choose to view an existing test or to start a new test. Because this is the first time, we will choose to configure and execute a new test using the parameters that we will specify in the next set of screens.

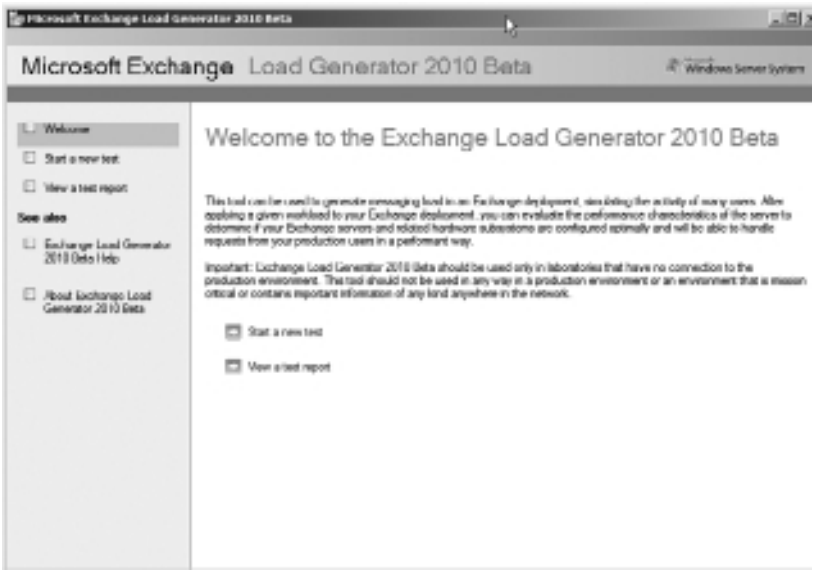
Microsoft Exchange Load Generator 2010 Beta
Welcome to the Exchange Load Generator 2010 Beta
Start a new test
View a test report

FIGURE 6.30
LoadGen 2010 Beta opening screen

After selecting the Start A New Test option, we are given the opportunity to define a new configuration or to use an existing configuration you might have saved in a file. Again, because this is a first-time run of the tool, select the Create A New Test Configuration option.

On the next screen, we must enter the parameters shown in Table 6.22 to enable the test.

TABLE 6.22: Exchange LoadGen 2010 Beta Table

LoadGen 2010 Parameter	Required User Input
Length Of The Simulation Or Workday	We selected 8 hours, which is usually sufficient.
Total Length Of Simulation	We selected 10 hours. We usually allow 2 hours for Exchange to reach steady state.
Server Stress Mode Option	Runs the simulation as fast it can. We did not select this option.
Stop Simulation After Number Of Exceptions	We did not select this option either.
Active Directory Forest Or Domain Controller Name	We entered the sourdisc domain in our case.
Account To Run The Simulation Under	We entered the Administrator account in this case.
Account Password	We entered password for the selected account.
Directory Access Domain	Again, we used the sourdisc domain.
Mailbox Account Master Password	We entered the password for the selected account.
Mailbox Account Master Domain	Again, we entered user the sourdisc domain.

Figure 6.31 shows the required user input screen and the values that we entered for our test simulation.

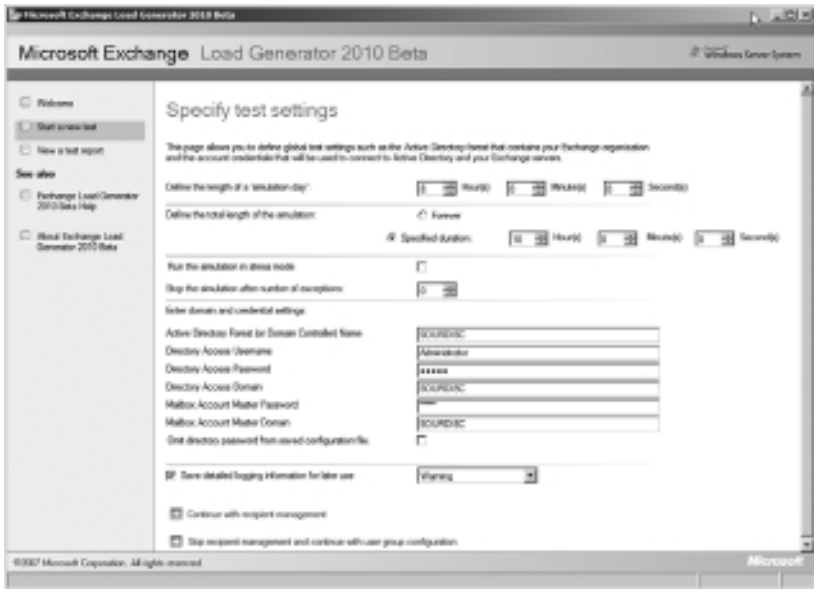

FIGURE 6.31
Specifying test settings screen

Once we have entered all the required input, select the Continue With Recipient Management option. Moving to the next screen, the topology is examined, and we can configure the number of users per detected databases.

Here we can tailor the distribution of your users per database. We can have some databases host more users with smaller mailboxes, while other databases host few users with larger mailboxes.

In Figure 6.32 we see the detected databases, and we enter the values for the number of users that we want to test per database. As we can see from the figure, we are evenly distributing the numbers of users per database.

On the next set of screens, we will get to enter the number of contacts and distribution lists, and finally we are offered the option of saving your configuration for future test simulation usage.

At this point, we initialize the test simulation configuration and then choose to run the simulation. Depending on the numbers of users, mailbox size, and user profile, the initialization can take quite a bit of time, so be patient!

After the test has completed, ideally without errors, we will be presented with a pass or fail summary screen. Unlike previous versions of the LoadGen tool that gave detailed information about the simulation, Microsoft has chosen to simplify the results of the simulation. Figure 6.33 displays the output of one of the simulations we ran for the proposed sourdisc.com Exchange 2010 Server deployment.

FIGURE 6.32
LoadGen 2010 Beta Exchange topology

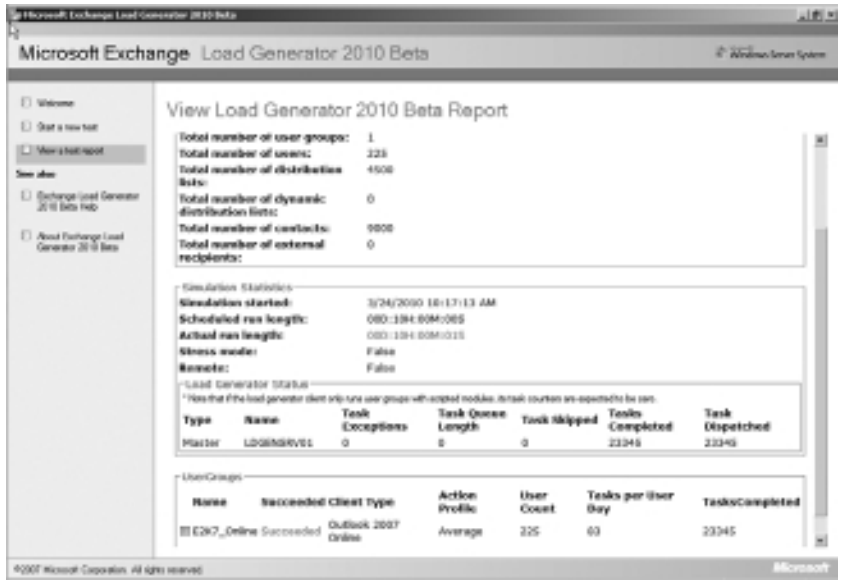

FIGURE 6.33
Exchange LoadGen 2010 Beta summary results

After numerous simulated runs of the proposed configuration with larger numbers of users, we have determined that the proposed configuration will more than support the number of users with

considerable available headroom. At this point, we can perform the following to prepare the virtual machines for deployment into the production sourdisc.com Active Directory forest:

1. Shut down the LoadGen 2010 Beta server.
2. Delete the databases on the Exchange Server 2010 virtual machines.
3. Uninstall Exchange Server 2010 from the virtual machines.
4. Place the virtual machines in workgroup mode.
5. Shut down the isolated sourdisc.com Active Directory forest.
6. Join the Exchange Server 2010 virtual machines to the production sourdisc.com Active Directory forest.
7. Reinstall all Exchange Server 2010 server roles on to the designated virtual machines.

At this point, we can see that the VIP Exchange Server 2003 cluster is still functional by opening the Cluster Administrator on the Exchange Server 2003 Mailbox server. Click the Start button on the desktop, and navigate to Administrative Tools ➢ Cluster Administrator to open the cluster administration tool.

In Figure 6.34, we display one of the Source Discovery clustered Mailbox servers, and we can see that the Windows Server 2003/Exchange Server cluster SDMEVS01 is still fully functional.

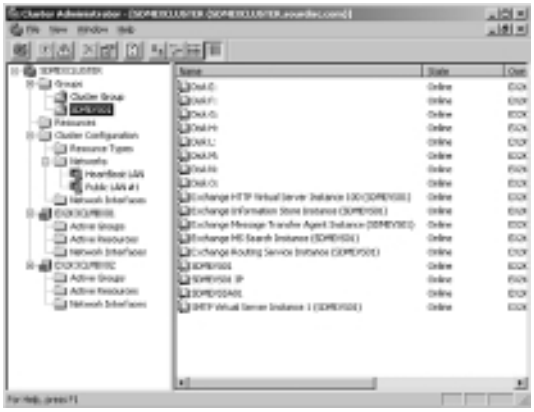

FIGURE 6.34
Windows/ Exchange Server 2003 cluster

With the Exchange Server 2010 virtual machines rejoined to the production sourdisc.com Active Directory forest and all the server roles installed on to the virtual machines, we can continue with your Exchange Server 2010 deployment.

LICENSING THE EXCHANGE SERVER 2010 SERVERS

As with any Microsoft operating system and application, you must license them to be in compliance with Microsoft end-user usage policies. Source Discovery has Microsoft volume licensing and premiere support agreements in place. We are deploying the Exchange Server 2010 messaging environment as well as all other applications and infrastructure services using the Microsoft Windows Server 2008 R2 Datacenter and Enterprise editions where applicable. Licensing the ESX hosts with Datacenter allows us to create and run as many virtual machines as you can on the ESX v4.0 hosts while maintaining licensing compliance with Microsoft.

For the Exchange Server 2010 messaging environment, we will be deploying Windows Server 2008 R2 Enterprise edition and Exchange Server 2010 standard licensing for the CAS and Hub Transport server roles.

Because Source Discovery has a population of users with very large mailboxes, you may be required to deploy more than five databases per virtual machine. In this case, you will utilize the Exchange Server 2010 Enterprise edition for the mailbox virtual machines.

TIP Microsoft provides the Exchange 2010 Mailbox Server Role Requirements Calculator to help you determine sizing requirements for your entire Exchange Server 2010 environment. You can download the calculator from the following site:

```
http://msexchangeteam.com/files/453145/download.aspx
```

Be sure to check back frequently because the calculator is updated often.

With the Exchange Server 2010 Enterprise edition licensing, we can have up to 100 databases per Exchange server host. Standard edition licensing allows up to five databases per host.

To license the Exchange Server 2010 virtual machines, we will open the Exchange Management Console (EMC) by clicking the Start button on the desktop. Select All Programs ➢ Microsoft Exchange Server 2010 Program Group ➢ Exchange Management Console. The EMC will open, and we can expand the Microsoft Exchange On-Premise container to reveal the Server Configuration Container.

Now expand the Server Configuration container, and in the Action pane choose Enter Product Key Group. Then enter the Exchange Server 2010 product key for the server that we are licensing.

TIP If you haven't entered licenses for any of your installed Exchange servers, you will get a pop-up window showing the servers that have yet to be licensed and the number of days left before you are forced to license them.

In the Source Discovery deployment, we have not licensed any of the Exchange Server 2010 virtual machines. Open the EMC, and as we can see in Figure 6.35, we get a pop-up reminder informing us that we need to license the Exchange Server 2010 server and the number of days left.

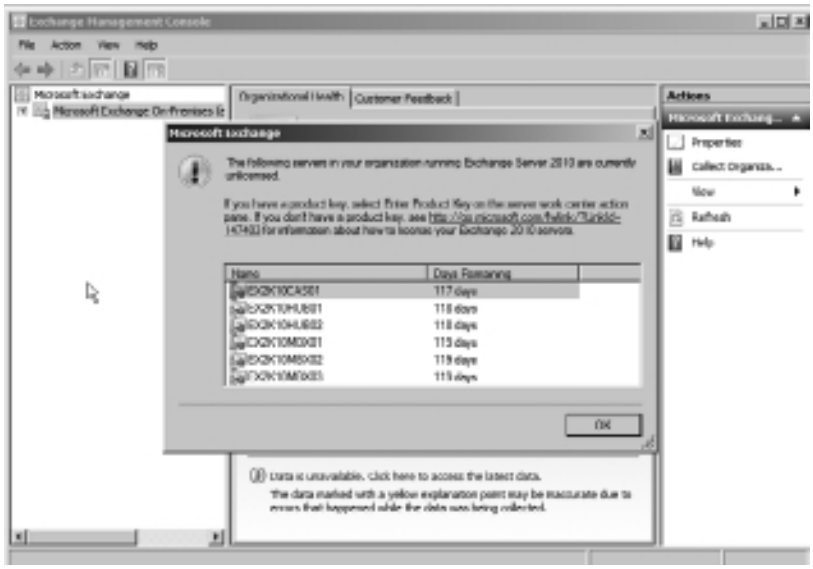

FIGURE 6.35 Exchange Server 2010 license notification

Click Enter Product Key in the Action pane; you are presented with the Enter Product Key Wizard, as shown in Figure 6.36. You can accomplish this through the Exchange Management Shell (EMS) as well. Once you have entered your product key for the Exchange Server 2010 server, you move through the remainder of the wizard's screens, and if you entered the correct key, you will receive a warning message informing you that the information store service will need to be restarted to enable the license.

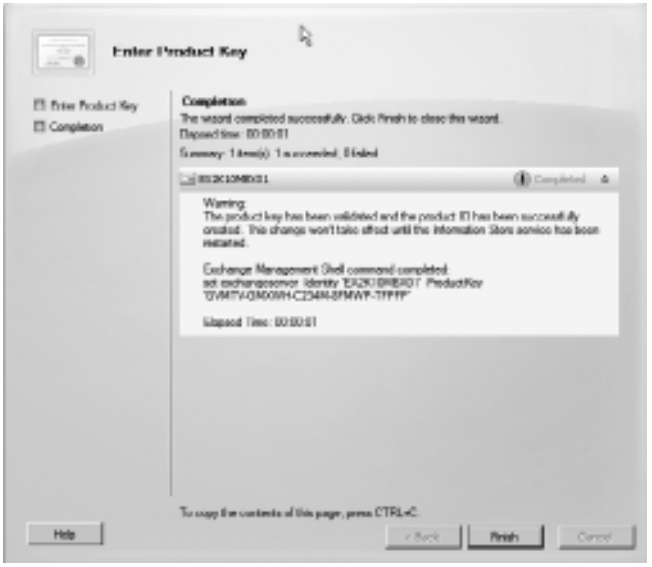

FIGURE 6.36
Entering the Exchange Server 2010 product key

Repeat this process for all the remaining Exchange Server 2010 virtual machines to be licensed, and after restarting the information store services on all Exchange Server 2010 virtual machines, you are now running licensed copies of the Exchange Server 2010 application in the Source Discovery messaging environment.

CREATING THE EXCHANGE SERVER 2010 DATABASE AVAILABILITY GROUP

Earlier we explained that we would be deploying Exchange Server 2010 DAGs for the executive, legal, and other staff VIPs in the Source Discovery organization. From here we will create a single DAG to hold the three Mailbox server role virtual machines.

DAGs can be created using the EMC or the Exchange Management Shell. We will create the Source Discovery DAG using a combination of both methods. When we create the DAG using the EMC, we cannot specify a static IP address for the DAG. This shortcoming will be addressed in the Exchange Server 2010 service pack 1 update.

A newly created DAG that is created from the EMC will use a DHCP address for the public-facing network communication and DAG NetBIOS network name. Although we can deploy DAGs with DHCP addresses, Source Discovery opted to assign a static IP address to its DAG.

Because of the static IP address requirement, we had to create the DAG and add members using both the EMS and the EMC. Although we could have done the entire process using the

EMS, we found it much easier to initially create the DAG using the EMS and add the group servers using the EMC.

Open the EMS by clicking the Start button and selecting All Programs ➢ Microsoft Exchange Server 2010 Program Group ➢ Exchange Management Shell.

From here the EMS will open, and you will see a PowerShell command window. From here, enter the following command at the command prompt:

```
C:\New-DatabaseAvailabilityGroup -Name SDEX2K10DAG01
-DatabaseAvailabilityGroupIpAddresses 192.168.100.220
```

This static address should be on the same subnet or routable to your public-facing network adapters. Figure 6.37 shows the output of the `Get-DatabaseAvailabilityGroup |FL` command. In the figure you can see the static IP4 address that we assign the DAG along with other important information about the DAG.

FIGURE 6.37
DAG properties

TIP You can find more information about Exchange MAPI networks and replication networks requirements on the Microsoft TechNet website for Exchange Server at the following URL:

`http://technet.microsoft.com/en-us/library.aa998636.aspx`

If you open the EMC and expand the Organization Configuration container, select the Mailbox object, and click the Database Availability Group tab, you will see that the new DAG has been created. Even though you have created a DAG, it is pretty much useless without anything to protect.

You must now add members to the DAG in order to create the highly available Exchange Server 2010 solution for Source Discovery. Open the EMC, and navigate to the Organization Configuration container.

Select the Database Availability Group tab, and click the Manage Database Availability Group Membership link under the DAG name. You will then be presented with the Manage Database Availability Group Membership Wizard.

From this point you select the Exchange Server 2010 server to add to the DAG. The server will be added to the DAG, and the network properties will be visible in the DAG Name Networks pane. You can expand the container objects in this pane for more information about the DAG network.

Figure 6.38 shows that we have configured all the Exchange Server 2010 virtual machines and they are now participating members of the SDEX2K10DAG01 DAG.

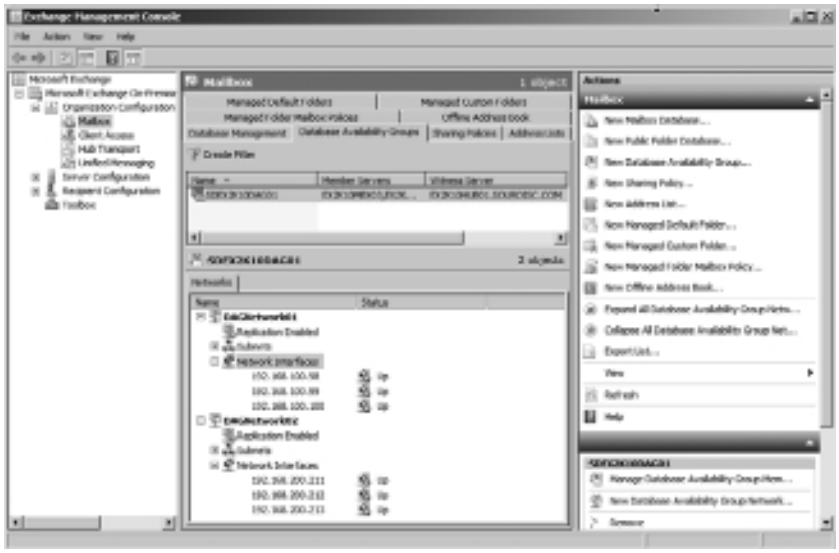

FIGURE 6.38
Showing the final DAG configuration

CREATING AN EXCHANGE SERVER 2010 DATABASE

We now come to the portion of the virtualized Exchange Server 2010 deployment where we must start to create the Exchange Server 2010 mailbox databases. An Exchange Server 2010 server database is now defined as the database and its transaction log files.

Exchange Server 2010 does away with storage groups. As we mentioned earlier in the chapter, each Exchange Server 2010 database is now an Active Directory object and must have a unique name.

This does not mean that you can create a bunch of databases and rename them. Each database must be unique in name and universally unique identifier (UUID). Failure to properly create your Exchange Server 2010 database can result in data corruption and the inability to even mount the databases at all.

You start creating the Exchange Server 2010 by opening the Exchange Server 2010 Management Console (EMC) and navigating from the Microsoft Exchange On-Premise container to the Organization Configuration container.

Expand this container, and select the Mailbox leaf of the container. In the Action area of the EMC, select the New Mailbox Database option. You are presented with the New Mailbox Database Wizard, as shown in Figure 6.39.

As you can see in the figure, you enter the name of the first Exchange Server 2010 mailbox database that you want to create. As you move through the screens, you will need to select which Mailbox server within the DAG will host the database.

In our DAG configuration, we configured three Exchange Server 2010 servers, so we have the choice of one of the three. We should mention that we are creating the databases and creating the copies of the databases while they are empty to minimize seeding time. Figure 6.40 shows the selections of the name and server to host the database.

FIGURE 6.39
New Mailbox Database Wizard

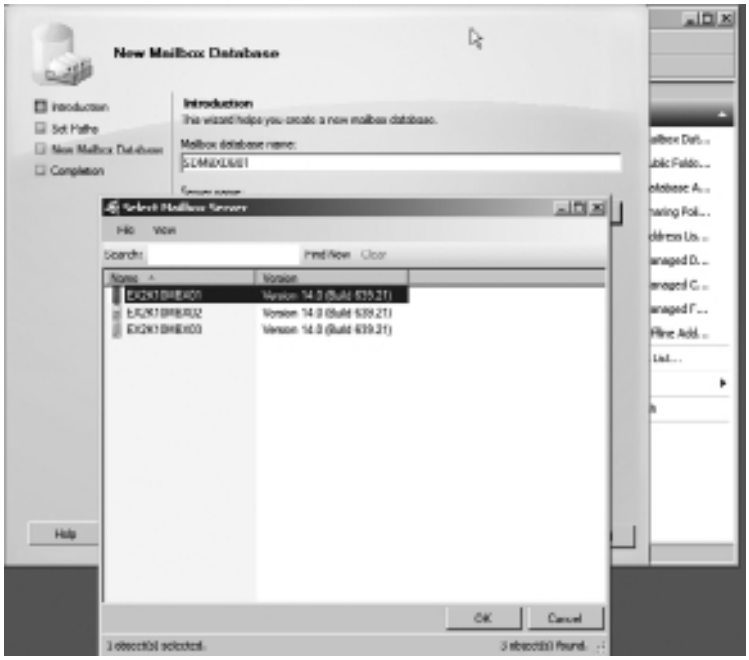

FIGURE 6.40
Continuing with the Exchange database creation

Moving into the next set of screens, you will be given the opportunity to choose where you want to create the database and its log files and whether to mount the newly created database. In this scenario, elect not to mount the databases, and create the database and its log files on the SAN volumes that you provisioned before deploying the virtual machines.

Figure 6.41 shows that the database creation was successful; click the Finish button to close the wizard.

FIGURE 6.41
Finalizing Exchange Server 2010 database creation

At this point, you will create the remaining databases and mount them so that you can start to perform mailbox moves from the Exchange Server 2003 messaging environment to the newly virtualized Exchange Server 2010 messaging environment.

CREATING ADDITIONAL EXCHANGE SERVER 2010 DATABASE COPIES

After the DAG is created and the Exchange Server 2010 virtual machines are added as members of the created DAG, you can start creating additional databases and create the Exchange Server 2010 database mobility solution.

This configuration will give the truly high availability that you are seeking. The active database will be in a mounted state, while the passive database copies can be any number of states to include healthy, suspended, and unhealthy, to name a few.

You can create additional database copies from the EMC as well as the EMS. For the sake of simplicity, create the database copies using the EMC. With EMC open, click the Database Management tab to reveal all the created databases.

You can select the database that you want to make a copy of, and in the Action pane under the database name, select the Add Mailbox Database Copy link to start the database copy creation process.

When the Add Mailbox Database Copy Wizard opens, select the server that you want to host the copy of the database and the Activation Preference number (this is the number that is used

by Exchange administrators to control the activation of Exchange database copies). Exchange will only activate healthy copies of the databases in the event of a database or server failure. Figure 6.42 displays the primary Exchange Server 2010 database and its copies.

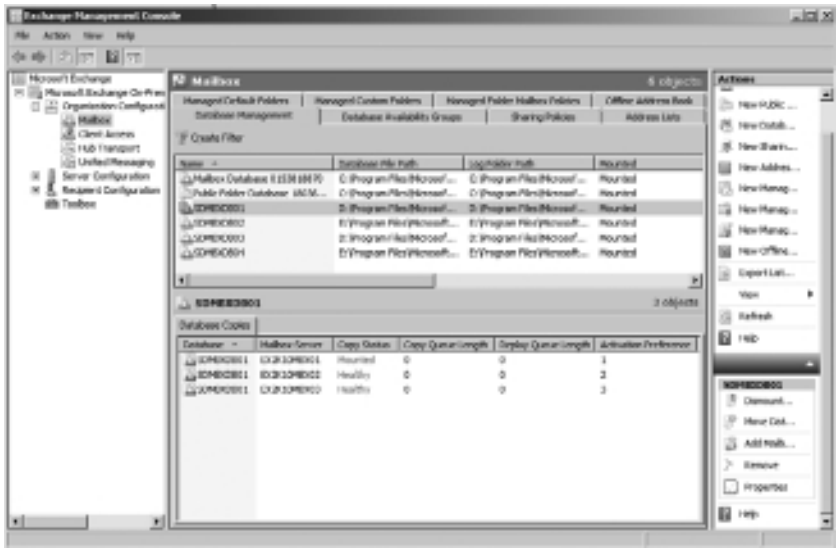

FIGURE 6.42
Exchange Server 2010 database copies

In Figure 6.42, you can see an Exchange Server 2010 database called SDMBXDB01 with two additional copies of that database that have been created. You will also notice that each database and its copy are hosted by a different Exchange Server 2010 DAG member, and the order of activation is numbered in accordance with how they are hosted by the DAG members. Be sure to allow ample time for the database copies to seed and the copy status to become healthy before moving any mailboxes to the Exchange Server 2010 messaging environment.

With all the Exchange Server 2010 databases created and their copies created as well, you are now ready to proceed with the movement of mailboxes from the Exchange Server 2003 messaging environment to the Exchange Server 2010 messaging environment.

Moving Exchange Server 2003 Mailboxes to Exchange Server 2010

Because the Source Discovery messaging environment consists of only Exchange Server 2003 servers, you can't perform online moves of mailboxes from the Exchange Server 2003. Source Discovery has designated maintenance schedules for performing the mailbox moves of the executives, lawyers, and other critical staff in the organization.

Again, you can move mailboxes from the Exchange Server 2003 messaging environment to the Exchange Server 2010 messaging environment using the EMS, but for the sake of simplicity, you decide to start the mailbox moves using the EMC.

WARNING Ensure that the Exchange Mailbox Replication Server on the CAS role has been started, or you will get error messages stating the Mailbox Replication Service is not running on any CAS server in the organization, and the mailbox move will fail.

Open the EMC, expand the Microsoft Exchange On-Premise container, expand the Recipient Container, and select the Mail leaf. In the Mailbox – Entire Forest pane, select the mailboxes that you want to move to Exchange Server 2010 messaging environment.

Figure 6.43 shows the selection of the mailbox users to move and the new database target for the mailboxes.

FIGURE 6.43
Selecting the users and database

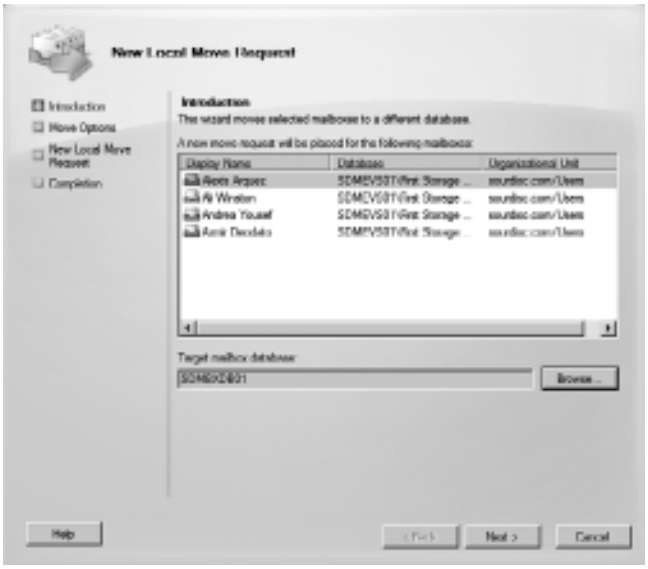

From here you click the Next button to specify the move options. You can specify the actions to take if and when you encounter corruption in the mailbox. You can opt to skip the mailbox if corruption is found. This will give you the chance to repair the mailbox before moving it into the Exchange Server 2010 environment.

From here you are presented with the summary screen. This is the last point at which you can choose to go back and correct or change options in the move request.

When you click the New button, the mailbox move will proceed, and if there are no errors or warnings, the wizard will complete, and you will proceed to the next set of users to move from the Exchange Server 2003 to Exchange Server 2010 messaging environment. Figure 6.44 displays the results of the mailbox move.

As you can see in the figure, all the selected users were moved without issue. From here you will move the remainder of the Source Discovery tier-1 users from the Exchange Server 2003 messaging environment to the Exchange Server 2010 messaging environment.

In keeping with Microsoft and VMware best practices, the Source Discovery Systems, VI, and Exchange administrators are all keeping a close watch on the newly deployed Exchange Server 2010 messaging environment. They have implemented a weekly full backup of the Exchange databases for insurance.

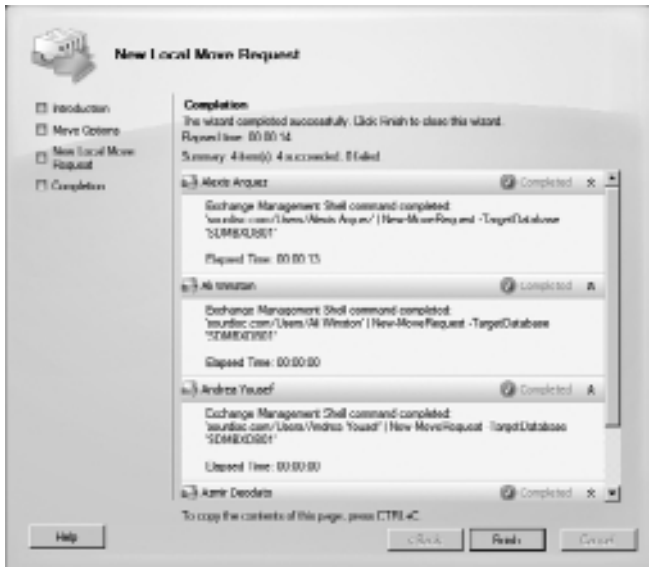

FIGURE 6.44
Move mailbox results

With the success of the newly deployed Exchange Server 2010 messaging environment, Source Discovery has proceeded with the move of the tier-2 onto stand-alone Exchange Server 2010 servers and is in the midst of moving the users to the new environment.

Summary

In this chapter, we covered some of the more compelling features of the Exchange Server 2010 product. Although there are a whole host of other features, we do not have enough room to cover all of them in any detail.

We covered briefly the Exchange Server 2010 server roles, their responsibilities, and their support in a virtualization environment. You learned that Unified Messaging is not supported virtualized.

We also talked about considerations for virtualizing the CAS, Hub, Edge Transport, and Mailbox server roles using vSphere. We went through a brief explanation of a true highly available solution and what it consists of. After the discussion of high availability, we discussed the parameters and business requirements for the Source Discovery proof of concept.

We manually prepared the Active Directory for Exchange Server 2010 deployment after preparing the Active Directory, and we shut down and cloned two Active Directory domain controllers. We showed how to build a test environment based on the live production and isolate it from the production network in order to perform the necessary testing to ensure the proposed solution would exceed the business requirements and service-level agreements.

After successful testing, we showed how to dismantle the test environment and build the live production environment. We licensed all the Exchange Server 2010 servers in the production Source Discovery messaging environment.

For the proof of concept, you created the production DAG and added the three Exchange Server 2010 mailbox virtual machines as members of the DAG. You also learned that DAGS created through the EMC are not given a static IP address and that you cannot assign a static IP address through the EMC.

This caused you to choose the EMS to create the DAG where you could assign an IP4 type address to the DAG NetBIOS network name. Source Discovery does not like to have critical business services dependent on DHCP services. Once the Exchange Server 2010 virtual machines were added as members of the DAG, you then created the primary databases and created their copies and distributed them among the DAG members. After allowing ample time for database copying and seeding to occur, you then proceeded to move the tier-1 users from the Exchange Server 2003 messaging environment to the Exchange Server 2010 DAG environment.

With the improvements of the Exchange Server 2010 product, reductions in storage I/O, more efficient use of hardware, and more powerful hardware, Exchange Server 2010 is more virtualization-friendly than ever. Remember when virtualizing any application to follow the recommended best practices for the design and implementation of that application.

VMware has published many test results for tier-1 applications that also include the Oracle database, SAP, and many health-care applications as well as Active Directory, SQL Server, Exchange Server, and SharePoint Server. vSphere is a mature virtualization platform that affords you a type of flexibility not found through the deployment of physical servers.

Chapter 7

Virtualizing SQL Server

Microsoft SQL Server is a widely deployed database server supported on Windows Server operating systems. SQL Server also serves as the most commonly used back-end database server for other Microsoft server applications based on IIS, including custom or homegrown applications as well as SharePoint. Microsoft continues to build SQL Server to be competitive with Oracle and IBM DB2, which are other popular database servers. Some industry analysts rank SQL Server as the number-one database based on the number of units deployed.

As virtualization becomes a standard practice within enterprises, many database administrators (DBAs) are embarking on SQL Server virtualization projects. SQL Server is a general-purpose database server that supports a wide variety of applications with very diverse resource requirements. You need to make several design and deployment decisions, such as specific application requirements (for example, performance, sizing, and availability) and virtualization capabilities (for example, scalability, licensing, support), when you deploy SQL Server in virtual environment.

This chapter will help SQL Server DBAs and IT administrators who want to virtualize SQL Server using VMware vSphere. It will address frequently asked questions and provide proof-of-concept methodology for SQL Server virtualization with VMware vSphere; it will also cover the considerations for virtualizing SQL Server on the vSphere platform and how to characterize existing SQL Server deployments for virtualization.

SQL Server Virtualization Considerations

As virtualization becomes the standard across IT data centers, CIOs and virtualization administrators are increasingly attempting to virtualize SQL Server deployments in their environments. DBAs have been skeptical about database server virtualization but have been warming up to wide-scale database virtualization in production environments recently. The database server is the last component in the application software stack that is responsible for persisting data to the physical storage, and hence DBAs have rightfully been conservative in adopting database virtualization.

We will discuss the following frequently asked questions regarding SQL Server virtualization:

- Is SQL Server data integrity maintained in virtual environments?
- How does virtualization compare to other SQL Server consolidation options?
- How does SQL Server perform when virtualized?
- Should we use 32-bit SQL Server or migrate to 64-bit version for virtualization?

- How are SQL Server high-availability and backup options affected in virtual environment?
- Should we dedicate a single ESX Server for all SQL Server virtual machines or combine these with other application virtual machines?
- Is it better to use fewer large SQL Server virtual machines or spread out SQL Server instances into a number of smaller virtual machines?
- What SQL Server databases should we *not* virtualize?

In the following sections, we will discuss these questions and more in detail.

Data Integrity in a Virtual Environment

One of the first questions DBAs often ask is whether database ACID properties and data integrity are maintained in a virtual environment, that is, whether the SQL Server deployments in virtual environments are I/O crash consistent. I/O crash consistency requires that a correct order of writes is maintained so that an application can restart properly in the event of a server crash. In other words, crash consistency makes sure that an I/O acknowledged successfully to the application is in fact stored persistently on disk. Because database is typically the last software component in an application stack, DBAs are on the hook to ensure that any application data is persisted and maintained correctly. Therefore, I/O crash consistency is an important consideration for databases, more so than other server applications.

VMware ESX or ESXi, which uses a bare-metal or hypervisor architecture, does guarantee I/O crash consistency. Consequently, SQL Server instances running inside Windows virtual machines on ESX have the same crash consistency guarantees as applications running on physical machines or physical disk controllers.

VMware-hosted products, such as VMware Workstation, Server, or Fusion, depend on the host operating system for I/O crash consistency. Deploying SQL Server virtual machines in hosted virtualization environments is not recommended and should be done only in test and development environments where data is not critical. We recommend you use VMware's bare-metal hypervisor—ESX or the free ESXi—to virtualize SQL Server even for test and development environments.

Figure 7.1 shows the difference between the hosted and hypervisor architectures used in virtualization products.

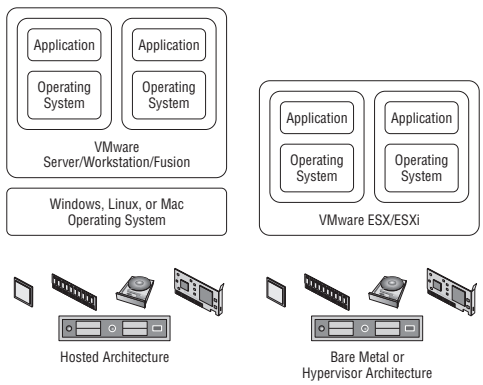

FIGURE 7.1
Hosted vs. hypervisor architecture

I/O Crash Consistency in VMware-Based Virtual Environments

VMware has discussed the I/O crash consistency in VMware-based virtual environments in the KB article 1008542. You can check the details at `http://kb.vmware.com/kb/1008542`.

SQL Server Consolidation Options

SQL Server provides its own options for consolidating databases (Figure 7.2 and Figure 7.3). These options include the following:

- Single SQL Server database with multiple schemas
- Single SQL Server instance with multiple databases
- Multiple SQL Server instances on single physical server

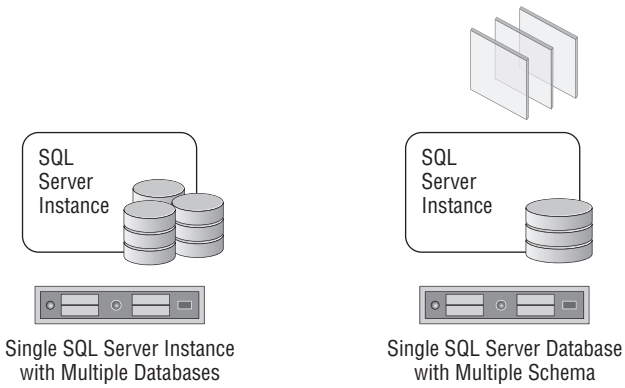
SQL Server Instance
SQL Server Instance
Single SQL Server Instance with Multiple Databases
Single SQL Server Database with Multiple Schema

Figure 7.2
SQL Server consolidation options using single SQL Server instance

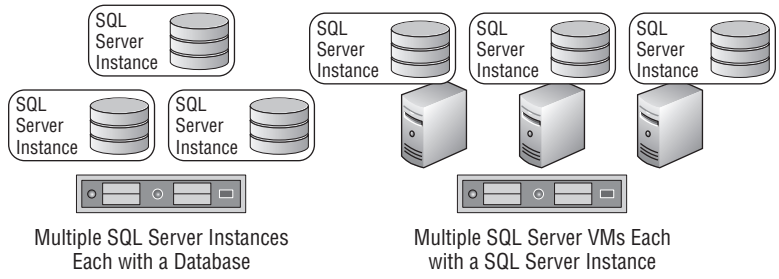
SQL Server Instance
SQL Server Instance
SQL Server Instance
SQL Server Instance
SQL Server Instance
SQL Server Instance
SQL Server Instance
Multiple SQL Server Instances Each with a Database
Multiple SQL Server VMs Each with a SQL Server Instance

Figure 7.3
SQL Server consolidation options: multiple SQL Server instances

With the introduction of 64-bit versions of Windows operating system and SQL Server, customers are being advised to consolidate SQL Server instances in a physical environment. Naturally, customers wonder how virtualized SQL Server deployments differ from the consolidation options in a physical environment. When multiple SQL Server instances are consolidated on a single physical server, they still share the same Windows operating system instance. In virtual deployment, however, each SQL Server instance resides inside its own Windows virtual machine. From a SQL Server

DBA's perspective, the key difference is that virtual SQL Server deployments offer complete isolation at various levels:

- Operating system
- Database workloads (OLTP/DSS)
- Security
- Performance/resources
- Configuration
- Fault isolation
- Change management

In virtual deployments, SQL Server DBAs can customize and lock down the SQL database configuration to meet the specific application- or department-level requirements. Furthermore, any changes made to the virtual deployment affect only the specific VM and thus affect only the users associated with it.

On the other hand, when consolidating SQL Server databases in physical environments, SQL Server DBAs often need to adopt a least common denominator configuration approach that is acceptable to all databases hosted on the server, and therefore they end up making various compromises during deployment.

SQL Server Performance in a Virtual Environment

Performance is the key to meeting any service-level agreements in production environments. This is especially important when virtualizing SQL Server given the perception that databases have too much I/O resulting in high performance overhead. VMware has published several performance studies using single and multiple VM configurations to demonstrate that virtualized SQL Server instances using VMware vSphere can deliver the performance to meet business requirements. You can find these VMware studies at `www.vmware.com/solutions/business-critical-apps/sql/resources.html`. There are a couple of key reasons why the vast majority of SQL Server deployments are good candidates for virtualization.

First, most SQL Server instances in physical environments use very little resources. The data published by VMware in its SQL Server best practices white paper, based on the analysis of its Capacity Planner data warehouse, states that average physical SQL Server instances

- use less than 6 percent CPU utilization,
- use approximately 60 percent of the average 3.1 GB of configured memory,
- perform about 20 disk I/O operations per second, and
- perform approximately 3 Mbps network I/O.

These smaller and lightly loaded SQL Server databases represent a large virtualization opportunity. Such databases have been regularly successfully migrated to virtual machines running on VMware ESX.

Second, virtualization technology has made significant advances in the past couple of years. Recent processors from AMD and Intel have hardware support for CPU instruction set and memory management unit (MMU) virtualization (Figure 7.4).

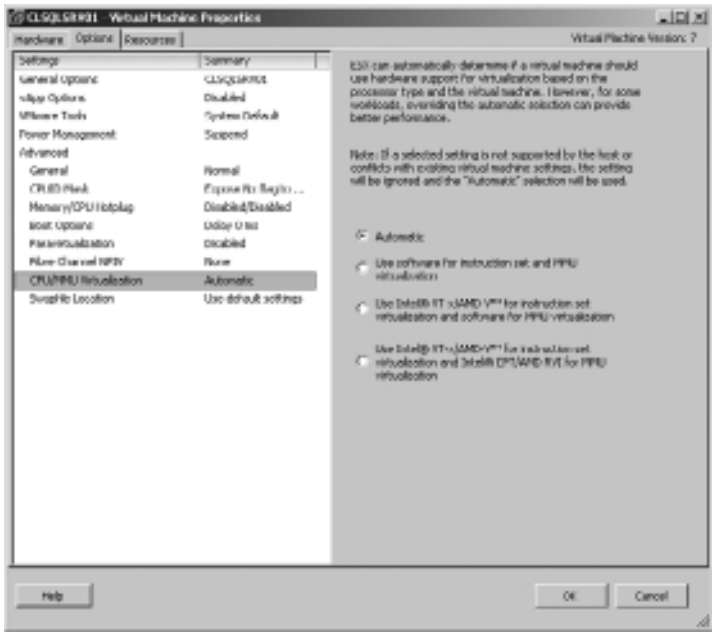

FIGURE 7.4
VM virtualization technology options

VMware vSphere also includes a number of storage improvements to address the performance concerns for I/O-intensive applications. Together these features significantly reduce the virtualization overhead and increase SQL Server performance as compared to native. As a result, any SQL Server deployment in a virtual environment using VMware vSphere and these newer processors will perform better compared to any proof of concepts you may have done, even within the past year.

32-bit vs. 64-bit SQL Server

As customers get ready to migrate to new hardware, OS, or SQL Server version, it is common for them to ask whether they should deploy the 32-bit or 64-bit version of SQL Server. We get the same question for SQL Server virtualization projects.

The 4 GB virtual memory limitations in the 32-bit environment are well known. Approaches to overcome the 4 GB virtual memory limit include using the 3 GB `boot.ini` switch and taking advantage of the Physical Address Extensions (PAE) in x86 processors. Both these approaches require recompiling the user applications. In both cases, the total virtual memory address space is still limited to 4 GB. In addition, when using the 3 GB switch, the kernel memory space is limited to 1 GB and may result in kernel memory exhaustion.

64-bit SQL Server does not incur the memory limits that the 32-bit version of SQL Server does. Hence, you hear about physical consolidation options such as multiple SQL Server databases or instances on a single, powerful 64-bit server. However, migrating to 64-bit SQL Server is not so

easy and may not offer all the benefits of a 64-bit environment for various reasons. The assumption that 64-bit will always outperform 32-bit is not always true and depends largely on the characteristics of the particular workload and hardware.

In physical environments, these factors can affect the SQL Server deployment decisions, especially when migrating to a 64-bit configuration. During SQL Server virtualization projects, you are not forced to make any immediate decisions. You can run both 32-bit and 64-bit SQL Server instances side by side in separate virtual machines on a single vSphere host. This flexibility implies no pressure to migrate to 64-bit environments or to port the legacy database applications.

NOTE Microsoft plans to ship 64-bit server operating systems from Windows Server 2008 R2. Windows Server 2008 is the last 32-bit server operating system from Microsoft. Also, installing a 32-bit SQL Server 2008 instance using WOW under 64-bit is not supported. This means either you cannot move certain SQL Server databases to 2008 or you have to port your applications, drivers, and other software to support 64-bit.

Licensing Considerations

We discussed licensing considerations in Chapter 2 for all Microsoft tier-1 applications covered in this book. We will reiterate the SQL Server–specific discussion here because it is eligible for per-processor licensing. SQL Server Enterprise belongs to a select group of Microsoft server applications, together with BizTalk Server Enterprise, Commerce Server Enterprise, and ISA Server Enterprise, which are eligible for per-processor licensing. For these products, you can run unlimited virtual instances if licensed for all physical processors in the server. Consequently, using per-processor licensing using SQL Server Enterprise edition may turn out to be quite favorable for your virtual deployment. Microsoft has also relaxed license reassignment restrictions, making tier-1 server application licensing more virtualization friendly.

Furthermore, SQL Server licensing is based on the number of processors, regardless of how many cores are in the processor. For example, if you are using SQL Server Enterprise edition on a four-processor server with quad-core processors, you will require only four processor licenses. When virtualizing SQL Server, customers often consolidate to fewer physical servers or migrate to newer multicore physical servers. In this process, you may end up requiring fewer processor licenses.

TIP For more information around licensing and volume licensing Microsoft applications and operating systems, please go to the Microsoft licensing website at `www.microsoft.com/licensing`. Also see "Application Server License Mobility" at `www.microsoft.com/licensing/resources/volbrief.mspx`.

Support Considerations

With databases being so critical to application data integrity, DBAs are concerned about SQL Server support in virtualization environments. SQL Server deployments in VMware environments are supported by Microsoft under the Server Virtualization Validation Program (SVVP). We discussed the SVVP in Chapter 2. Currently VMware ESX/ESXi 3.5 Updates 2 through 5 and vSphere 4.0 and 4.0 Update 1 are validated and supported under the SVVP. For more information on the Microsoft Server Virtualization Validation Program, see the Microsoft Windows Server Catalog website at `www.windowsservercatalog.com`.

In the next few steps, we will walk you through the SVVP Support Policy Wizard to verify SQL Server 2008 support when running inside Windows Server 2008 virtual machines on VMware vSphere 4.0. On the SVVP web page, click Support Policy Wizard. Start by selecting SQL Server 2008 from the product's drop-down list (Figure 7.5).

FIGURE 7.5
SVVP Support Policy Wizard, selecting a product

Next, select VMware vSphere 4 as the virtualization platform, and select 64-bit Windows Server 2008 as the guest operating system (Figure 7.6).

FIGURE 7.6
SVVP Support Policy Wizard, selecting virtualization platform and guest OS

The Support Policy Wizard shows that all versions of SQL Server 2008 running inside 64-bit Windows Server 2008 virtual machines on VMware vSphere are supported by Microsoft (Figure 7.7).

FIGURE 7.7 SVVP Support Policy Wizard, summary support statement

Another concern DBAs have is whether SQL Server failover clusters in virtual environments are supported by Microsoft. The support statement displayed by the SVVP Support Policy Wizard also clarifies that support. It explicitly states that Guest Failover Clustering is supported for SQL Server 2005 and SQL Server 2008 in a virtual machine for SVVP-certified configurations provided the requirements outlined in Microsoft Knowledge Base article 956893 are met. You can find more information about KB article 956893 at `http://support.microsoft.com/?id=956893`. With VMware vSphere, you can meet all the requirements listed in that KB article and deploy supported SQL Server failover clusters in your virtual deployment.

For non-SVVP-validated configurations of ESX and Microsoft server applications, support options vary depending on how customers purchase VMware and Microsoft products. For customers with a valid Microsoft Premier-level support agreement, Microsoft will use "commercially reasonable efforts" to support its products running on VMware virtual machines. If you have purchased VMware products together with server hardware from server resellers such as Dell, Fujitsu, Fujitsu-Siemens, HP, IBM, and Unisys, you can avail the end-to-end support for Microsoft software as part of a valid support agreement with the server reseller. This provides customers with one-stop support via the server reseller if an issue arises.

Security Considerations

Data security is a key requirement for any database deployment. DBAs often ask us whether they have to make any security compromises when virtualizing SQL Server. On the contrary, virtualizing SQL Server allows for more granular security configuration compared to native SQL Server consolidation. You can minimize the SQL Server "surface area" per virtual machine by customizing security at the VM level, installing and configuring components and features only as needed, and enabling data access protocols as needed. The database administrators can be more selective when assigning administrative and database owner privileges. Furthermore, separate virtual machines also make SQL Server patching and change management less risky for database and system administrators.

Hardware Considerations

Virtualization technology has changed significantly in the past few years. Both AMD and Intel have embraced virtualization and developed new processor features to overcome the initial shortcomings in x86 architecture. Newer processors from AMD and Intel now include hardware support for both CPU instruction set and memory management unit (MMU) virtualization. With these enhancements, hardware-assisted virtualization is now faster compared to the binary translation technique that VMware invented.

When considering hardware for your virtualization projects, we recommend you use the newer generation of processors from AMD and Intel. Since the virtualization is offloaded to hardware, virtualized applications perform much better with lesser virtualization overhead. In fact, with vSphere VMware prefers hardware-assisted virtualization over binary translation when creating virtual machines. VMware ESX 4.0 introduces support for the Intel Xeon processors that support EPT. Support for AMD RVI has existed since ESX 3.5. With the new hardware and vSphere, make sure to use the latest version of virtual hardware so you can leverage the latest scalability features such as eight virtual CPU support and better optimized virtual I/O drivers such as VMXNET3 and PVSCSI.

Storage Considerations

Storage configuration is critical to any successful database deployment. As with any physical SQL Server deployment, you should size the storage in virtualized environments properly to meet the database I/O requirements. Microsoft offers predeployment storage best practices information to help you configure storage for SQL Server environments.

TIP You can find SQL Server best practices at the Microsoft SQL Server TechNet page:

`http://technet.microsoft.com/en-us/sqlserver/bb671430.aspx`

You can find the Pre-deployment I/O Best Practices Article at `www.microsoft.com/technet/prodtechnol/sql/bestpractice/pdpliobp.mspx`.

Follow these guidelines for deployment on VMware vSphere as well.

There is one key difference, however. In physical SQL Server deployment, the data, transaction log, and TempDB LUNs are typically dedicated to a single host and sized accordingly. In virtual environments, you will be running multiple virtual machines, which are very likely to consolidate many different SQL Server workloads onto a single ESX host. Your storage subsystem should provide sufficient I/O throughput as well as storage capacity to accommodate the total needs of all virtual machines running on the ESX host. Most SQL Server performance issues in virtual environments result from such improper storage configuration.

Multiple heavily used virtual machines concurrently accessing the same VMFS volume, or multiple VMFS volumes backed by the same LUNs, can cause decreased storage performance. You must configure storage architecture appropriately to avoid this issue. We recommend you separate the data, transaction log, TempDB, and SQL Server partitions to different LUNs, if possible, on separate physical disk drives for each SQL Server virtual machine. If architected correctly, a shared storage model strategy used in virtualized deployments will provide many benefits, such as more effective storage resource utilization, reduced storage whitespace, better provisioning, and mobility using VMware VMotion and VMware Storage VMotion.

Shared Storage Scalability in Virtual Environment

A typical VMware ESX deployment includes multiple VMware ESX hosts sharing the same storage. For example, a LUN may be shared by multiple VMs and presented to multiple ESX hosts. Such sharing can impact the storage I/O performance because of several factors such as I/O queuing, SCSI reservations, and storage multipathing. VMware has discussed these issues and published results from their performance studies in white paper published at `www.vmware.com/files/pdf/scalable_storage_performance.pdf`.

In addition, when you design the storage for a SQL Server virtual environment, customers often ask us questions about which they should use:

- Fibre Channel or iSCSI storage protocol
- VMware Virtual Machine File System (VMFS) or raw device mapping (RDM) disks

Fibre Channel or iSCSI Storage Protocol

Customers usually do not choose between Fibre Channel and Ethernet-based storage protocols such as iSCSI or NFS simply for virtual SQL Server deployment. They usually make this decision based on in-house skills, cost, or manageability concerns. The storage infrastructure is typically already in place and is leveraged for virtualization projects. In case of either protocols, the storage throughput that vSphere can scale to is limited only by the link speeds. If you have made investments in the latest I/O technologies, be assured that vSphere supports the latest technologies such as 8 Gbps Fibre Channel connectivity, Fibre Channel over Ethernet (FCOE), and 10 GigE connectivity for IP-based storage.

Storage Protocol and Performance in vSphere Environment

VMware has published performance studies that compare storage I/O performance with different storage protocols such as Fibre Channel, hardware iSCSI, software iSCSI, and NFS. The results are posted in a white paper at `www.vmware.com/files/pdf/perf_vsphere_storage_protocols.pdf`.

VMFS or RDM

When you configure your VMware vSphere environment, you must decide whether to use RDM or VMFS partitions to store your virtual machines (Figure 7.8).

Using VMFS for most virtual disk storage is recommended because VMFS provides much more flexibility and in almost all cases performs just as well as RDMs. VMware reports very little I/O throughput difference between RDM and VMFS configurations. You should decide which storage configuration to use based on other aspects of your deployment. For example, you can move virtual machines using VMFS more easily than those using RDM. Using templates and self-provisioning is more complicated when you use RDM. Using RDM requires more storage administration because you need to provision separate LUNs for each virtual disk. RDM may be desirable in a small

number of cases, such as when you use MSCS-based failover clustering for SQL Server or when you leverage backup solutions that leverage SAN features such as snapshots. Using RDM can also help you during the initial transition to a VMware vSphere environment. However, after you complete the migration to a virtual environment, we recommend moving to VMFS for easier storage administration.

FIGURE 7.8
Virtual machine disk options

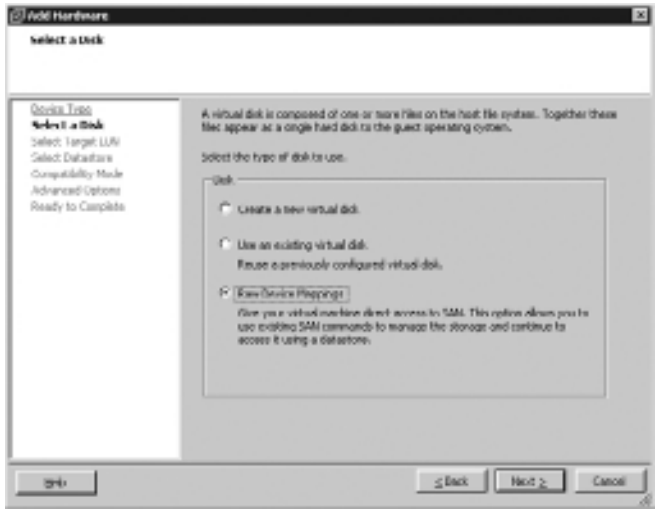
Add Hardware
Select a Disk
Device Type
Select a Disk
Select Target LUN
Select Datastore
Compatibility Mode
Advanced Options
Ready to Complete
A virtual disk is composed of one or more files on the host file system. Together these files appear as a single hard disk to the guest operating system.
Select the type of disk to use.
Disk
Create a new virtual disk
Use an existing virtual disk
Reuse a previously configured virtual disk.
Raw Device Mappings
Give your virtual machine direct access to SAN. This option allows you to use existing SAN commands to manage the storage and continue to access it using a datastore.
< Back
Next >
Cancel

VMFS and RDM Performance in vSphere Environment

VMware has published performance studies comparing VMFS and RDM disk partitions. The results described in these studies show that VMFS and RDM provide similar I/O throughput. For ESX 3.5, the results are published in the whitepaper `www.vmware.com/files/pdf/performance_char_vmfs_rdm.pdf`.

High Availability Considerations

Enterprises have various high-availability requirements for database deployments, which determine the amount of downtime allowed in a particular deployment. These requirements are specified in a number of nines, ranging from a single nine (90 percent availability, or 21.6 hours uptime daily) to five or more nines (99.999 percent availability, or approximately 5 minutes and 15 seconds per year downtime). SQL Server provides a range of high-availability features to help organizations meet these requirements. These features include log shipping, database mirroring with various levels of protection, and failover clustering using Microsoft Cluster Services (MSCS). Table 7.1 summarizes the level of high availability these features provide.

TABLE 7.1: Levels of High Availability Provided by Key SQL Server Features

TECHNOLOGY	GRANULARITY	DATA LOSS	DOWNTIME
Failover clustering	Server	No	~30 sec + recovery
Database mirroring	Server	No	< 3 secs
High safety mode with automatic failover	Database		
Log shipping	Database	Possible transaction log	Variable
Database mirroring	Database	None	< 3 secs
High-protection configuration			
Database mirroring	Database	Some	< 3 secs
High-performance configuration			
Replication	Publication	Some	Variable

Naturally, DBAs want to know whether virtualization changes any of the SQL Server high-availability options. You can continue to leverage the common SQL Server high-availability implementations such as failover clustering, database mirroring, and log shipping with VMware vSphere. In addition, VMware vSphere includes its own high-availability features to address both planned and unplanned downtimes caused by failures ranging from component problems to data center issues. VMware Distributed Resource Scheduler (DRS), VMware High Availability (HA), VMware Fault Tolerance (FT), and the VMware Site Recovery Manager (SRM) make it possible to eliminate planned downtime, reduce unplanned downtime, and recover rapidly from component, systematic, or site-wide outages.

Mission-critical databases that require three to five nines of availability (well under 10 percent) typically use failover clustering solutions such as Microsoft Clustering Services. However, because of the costs and complexity involved in implementing these solutions, other databases either implement SQL Server–provided high-availability options or have no high-availability solution. VMware HA provides an out-of-the-box high-availability solution for SQL Server databases and has proven to be more than sufficient for up to 99.99 percent high availability. We recommend using VMware HA for SQL Server databases that are not configured with any availability options. Together with SQL Server database mirroring and VMware HA, you can achieve a SQL Server "always on" strategy with the database operating in an unprotected state for minimal time.

In Chapter 3, we discussed that VMware vSphere offers several possible failover clustering configurations:

Cluster in a box (CIB) CIB is a Microsoft cluster in which all nodes of the cluster reside on the same ESX host and that uses `.vmdk` disks as shared storage. All networking is handled through the single ESX host.

Cluster across boxes (CAB) A Microsoft cluster places the nodes of the cluster on different ESX hosts. Networking is handled through dedicated vSwitches on each ESX host. Storage is provided using Raw Device Mappings (RDM) disks mounted on a virtual disk controller that is configured to share their disks across ESX hosts.

Clustering physical machines with virtual machines In this configuration, a virtual machine cluster node is added to the physical cluster. The networking is handled through the vSwitch on the ESX host that the clustered virtual machine is residing on. Storage is handled in the same way as in the CAB configuration in which RDMs are used to provide the storage used in the physical cluster to the virtual cluster node.

Although the CAB configuration provides both hardware- and software-level protection by placing the SQL Server cluster nodes on separate ESX hosts and is recommended for production use, the CIB configuration can be very helpful in test and development environments.

Virtualization also greatly simplifies disaster recovery between production and failover sites. As mentioned earlier, any disaster recovery solutions based on SQL Server high-availability features such as log shipping and database mirroring should work as usual in vSphere deployment. Other disaster recovery solutions that are application agnostic often require you to maintain identical server hardware at the disaster recovery site to prevent issues with hardware and device driver compatibility when servers are brought online at the recovery site. VMware vSphere using the SRM plug-in eliminates this requirement, and you can quickly bring systems online on any supported ESX host machines running anywhere in the world. Each virtual machine is captured in a distinct set of files, so replication to a disaster recovery site is simple, and systems can be back up and running quickly. You can leverage the SAN replication features provided by storage vendors for this VM copying.

Backup Considerations

In a vSphere environment, you have a wide range of backup/restore options. These options include the following:

- Use backup tools inside virtual machines.
- Use backup agents inside the virtual machines.
- Use backup agents in the ESX service console.
- Use VMware Data Recovery to perform a server-level backup.

Review the backup and restore options for Windows virtual machines discussed in Chapter 3. When you virtualize SQL Server, you can leverage the same options. vSphere supports the Volume Shadow Copy Service (VSS) framework, so you can use a VSS-compliant solution to create application-consistent copies of data with any of these backup options.

For smaller environments, you can use Windows Server Backup as an in-guest backup solution. You can also use VMware Consolidated Backup (VCB), VMware Data Recovery (VDR) tool, or the snapshot capabilities provided by your storage vendor.

This means that you can be very flexible when addressing your SQL Server database backup needs. Depending upon the recovery point objectives and recovery time objectives for your SQL Server databases, you can define which of the previous options is appropriate for your environment. You can continue to leverage the existing backup mechanism, or you can migrate to a cost-effective backup solution that provides a higher degree of flexibility and functionality than they currently use.

Although your backup and restore tools do not change for virtualized instances of SQL Server, you need to plan carefully to establish time windows for backing up and restoring data for different databases. Because backups generate an additional I/O load on the SQL Server virtual machine, you need to create a backup strategy that conforms to off-hours duration and that will adhere to existing service-level agreements. Multiple SQL Server databases can be consolidated on a single ESX host, so you will also need to stagger your backup and restore tasks so that the I/O subsystem can handle the load.

What Not to Virtualize

vSphere supports virtual machines with up to eight virtual processors and 255 GB of memory. A single vSphere host has been demonstrated to scale up to 350,000 disk I/O operations per second, and possibly more depending upon storage available. These represent the high watermark capabilities of VMware vSphere. Any SQL Server databases that use resources within these limits can be considered for virtualization using VMware vSphere. If a SQL Server database uses resources beyond these vSphere capabilities, you should consider scaling out the deployment so that applications requirements can be met. You should check whether the SQL Server instance or database can be scaled out to separate instances or databases onto more virtual machines or ESX hosts.

The key factor is the actual resource usage—CPU or memory utilization or disk I/O throughput—not simply the capacity of the physical hardware deployed for physical SQL Server. Make sure you normalize the CPU utilization or disk I/O throughput to the hardware and storage you plan to use for ESX. Whether a SQL database can be virtualized will also depend on other such requirements as capacity headroom, availability considerations, and so on.

In general, the combination of vSphere performance enhancements and hardware-assisted virtualization in AMD and Intel processors has raised the performance for any virtualized enterprise application. Any virtualization proof of concept or deployment projects using this combination will demonstrate performance improvements compared any virtualization attempts you may tried earlier.

Characterizing SQL Server Deployments for Virtualization

To evaluate existing SQL Server deployments as candidates for virtualization, you need a clear understanding of your organization's business and technical requirements for each SQL Server instance. These requirements span several dimensions, such as availability, performance, scalability, growth and headroom, patching, and backups.

When you consider virtualizing SQL Server deployments, the following are two key concerns:

- How to determine which SQL Server instances are good virtualization candidates
- What processes to use when migrating to a virtual environment

Follow these steps to simplify the process of identifying SQL Server candidates for virtualization and ensure successful deployment:

1. Characterize SQL Server workloads.
2. Define availability requirements.

3. Define backup strategies and requirements.
4. Define patching and update strategies.
5. Capture a resource utilization baseline for existing databases.
6. Define system capacity and growth requirements.
7. Rank SQL Server candidates for virtualization.
8. Plan a migration strategy.

Characterizing SQL Server Workloads

Most database workloads fall into two major categories: online transaction processing (OLTP) and decision support system (DSS) workloads or relational data warehouse (RDW) workloads.

Online transaction processing OLTP workloads consist of a high frequency of short transactions using a limited number of simple, predefined queries. OLTP queries usually affect small amounts of data and are expected to respond within a few seconds. OLTP workloads are typically used by a large number of users simultaneously, and transaction isolation and locking are key to support the high concurrency. OLTP workloads generally result in random reads and writes against data files. The data file read activity tends to be constant, while the data file write activity occurs during checkpoint operations. Both log reads and writes are sequential in nature and may vary in size.

Decision support system DSS workloads, on the other hand, consist of long-running, complex queries that can be either ad hoc or programmed using special online analytical processing tools. DSS queries usually affect large amounts of data (typically several terabytes in size) and tend to run much longer (often several minutes or even few hours). As a result, DSS queries are highly resource intensive and can saturate CPUs, memory, and storage subsystems. For these workloads, both reads and writes tend to be sequential in nature because of table or index scans and bulk insert operations. The I/O size is generally any multiple of the default 8 KB SQL Server page size, and read-aheads are commonly used.

These two workload types often present conflicting requirements for database administrators. Most DBAs either end up provisioning several separate SQL Server instances or compromise on the SQL Server configuration parameters when consolidating on a single SQL Server instance.

For example, when running a SQL Server instance as both an OLTP system and a BI reporting system, the OLTP users often suffer because of long-running reports. A long-running DSS query, which is typically processed in parallel streams on multiple processors, can easily consume all of the system resources, thus starving out other processes. To configure SQL Server to allow more resources for the online users, DBAs have two options. DBAs can permanently disable the allowed degree of parallelism by setting the "max degree of parallelism" configuration option. Alternatively, they can set this option to 1 during daytime hours when online users are active, and set it to 2 or higher after hours when batches are run that could take advantage of parallelism. Virtualizing SQL Server instances alleviates this; the SQL Server

DBAs can provision separate VMs for such workloads and optimize the SQL Server configuration as needed.

Capturing Resource Utilization Baseline for Existing Databases

After you clearly understand your organization's needs—business and technical requirements, availability, and other operational requirements for implementing SQL Server— the next important step is to establish a baseline using data from the currently running physical deployment. You can accomplish this in multiple ways:

- *You can use the virtualization assessment exercise delivered by VMware or its partners.* Using VMware Capacity Planner, this exercise collects all essential performance metrics—processor, disk, memory, and network statistics, including specific SQL Server metrics—on existing database servers in the environment. Capacity Planner analyzes this data to recommend server consolidation opportunities.
- *You can monitor physical SQL Server deployments to capture the baseline information yourself.* To help you in this process, we have provided you with the list of SQL Server metrics and other essential system metrics. You can use Windows Perfmon, SQL Server system tables, stored procedures, and dynamic management views to capture this information. For SQL Server 2008, you can also leverage the performance data warehouse to collect this information.

It is important to collect the data over a period of time long enough to reflect all variations in the usage patterns of SQL Server in your organization. The duration can range from one week to one month, depending on seasonal or cyclical usage peaks—when they occur, their intensity, how long they last, and other factors. This exercise not only helps you understand what resources your current physical SQL Servers are using, but it also makes it easier for you to create a short list of SQL Server instances to virtualize and to determine the order in which you should virtualize the SQL Server instances.

The Capacity Planner tool collects up to 30 days of system and performance data. The tool then analyzes the data to recommend which servers can be virtualized. Capacity Planner can also model various server consolidation scenarios based on inputs such as hardware configuration, processor models, and so on. The reports also include a preliminary total cost of ownership analysis of the sampled environment and recommended virtual deployment. You can access all this information from the web-based Capacity Planner Dashboard (Figure 7.9).

VMware Capacity Planner services are provided by both VMware Professional services that follow the VMware methodology for virtualization as well as VMware partners. Visit the VMware Capacity Planner website at the following URL for more information: `www.vmware.com/products/capacity-planner/`.

If you plan to monitor the physical SQL Server instances yourself, Table 7.2 describes the performance counters you can capture to get baseline resource utilization in a physical environment.

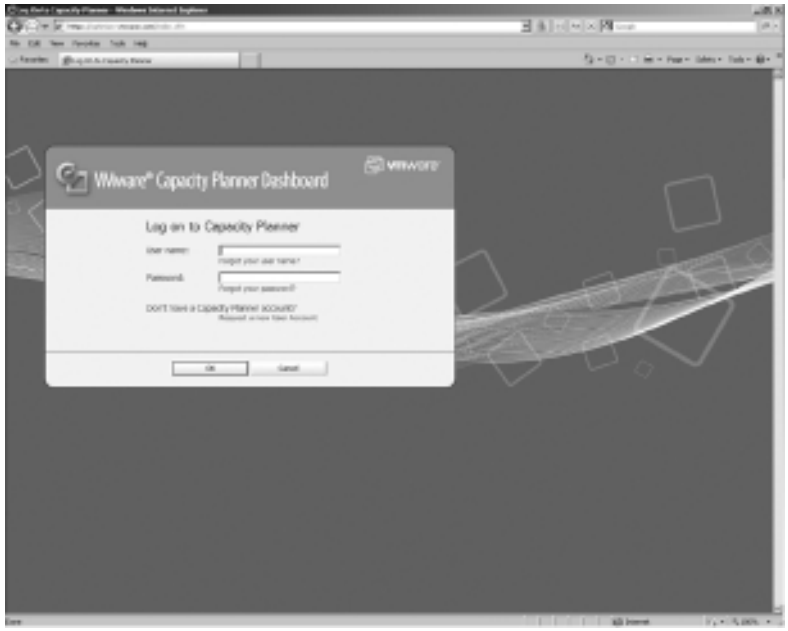

FIGURE 7.9
VMware Capacity Planner Dashboard

TABLE 7.2: Performance Counters to Capture

PERFORMANCE OBJECT	COUNTER	NOTES
Process	% Processor Time % Privileged Time Private Bytes Virtual Bytes Page Faults/sec Thread count Working Set I/O Read Bytes I/O Read Operations I/O Write Bytes I/O Write Operations	Information about the processes running on the server; check the SQL Server processes.
Processor	% Processor Time % Privileged Time	Total and per individual processor (especially with multiple cores); shows the performance of the processors.
Memory	Pages/sec Available Mbytes	Shows whether memory is paging to disk. Displays the amount of memory.

TABLE 7.2: Performance Counters to Capture *(CONTINUED)*

PERFORMANCE OBJECT	COUNTER	NOTES
Network Interface	Packets received/sec Packets sent/sec Bytes received/sec Bytes sent/sec	Per each network card, especially ones being used for SQL Server.
Physical Disk	%Disk Read Time %Disk Write Time %Idle time Avg. Disk Bytes/Read Avg. Disk Bytes/Write Avg. Disk Read Queue Length Avg. Disk Write Queue Length Avg. Disk sec/Read Avg. Disk sec/Write Split I/O/sec	Configure per disk as well as _Total; will show how I/O looks at a physical level for the disk configuration.
SQL Server: Buffer Manager	Buffer cache hit ratio Lazy writes/sec Read-ahead pages/sec Total pages	Percentage of pages found in the buffer pool, thus saving a read from disk. Number of buffers written by buffer manager's lazy writer. Number pages SQL Server reads in advance. Number of pages in the buffer pool (includes database free and stolen).
SQL Server: Access Methods	Full scans/sec Index searches/sec	Number of unrestricted full scans. Can be either base table or full index scans. Number of index searches.
SQL Server: Databases	Transactions/sec	Select per database; tells how many transactions are happening in the database at that given time.
SQL Server: General Statistics	User Connections Active Temp Tables	Gets the total amount of user connections in SQL Server. Indication of TempDB usage (but not necessarily the only indicator).
SQL Server: Memory Manager	Connection Memory(KB) Total Server Memory (KB)	Amount of memory consumed by connections. Amount of dynamic memory being used by the instance of SQL Server.

TABLE 7.2: Performance Counters to Capture *(CONTINUED)*

PERFORMANCE OBJECT	COUNTER	NOTES
SQL Server: SQL Statistics	Batch requests/sec SQL compilations/sec SQL recompilations/sec	Number of SQL batch requests received by server. Number of SQL compilations. Number of SQL recompilations.
SQL Server:Locks	Average wait time (ms) Number of deadlocks/sec	Average amount of wait time (milliseconds) for each lock request that resulted in a wait. Number of lock requests that resulted in a deadlock.

Defining Availability Requirements

When planning to virtualize SQL database servers, define your availability requirements to meet service-level agreements and handle peak load demands. Your deployment solution needs to address the applicable uptime guarantees, accounting for both planned and unplanned downtimes for various system and site components. Here are some of the things to consider:

- How much uptime guarantee do you need? How do you deal with planned downtime? How much unplanned downtime can you handle per application? Do you need a DR plan/site?
- Can you map capabilities to requirements, such as recovery time objectives (RTO) vs. recovery point objectives (RPO), VMware HA/DRS, clustering, database mirroring, log shipping, and so on?

Understanding these availability requirements helps you decide how to leverage such advanced features of VMware vSphere as MSCS support, DRS, HA, and SRM for your SQL Server virtual machines, together with or instead of SQL Server high-availability features.

Designing the proper availability strategy plays a large part in whether you succeed or fail in your SQL Server consolidation. After designing and deploying new virtualized SQL Server solutions, test your HA scenario to ensure that it meets or exceeds you service-level requirements.

Defining Backup Strategies and Requirements

A well-thought-out backup and recovery strategy and procedure can enhance your business operations around SQL Server availability and recovery. This can greatly shorten recovery times and help ensure that you meet your service-level agreements.

Consider the amount of data to be included in the backup, frequency of the backups, and time to restore that data in case of SQL Server VM failure. VSS-enabled snapshots can provide a rapid restore mechanism for SQL Server. Consult your storage/software vendor for available solutions that might meet your requirements.

When defining any backup and restore time windows, keep in mind that multiple databases will be consolidated on a single vSphere server.

Defining Patching and Update Strategies

Database maintenance and OS patching should play an important role in your consolidation strategy. Try to place instances that share a common maintenance schedule on the same ESX servers. VMware VMotion allows you to relocate virtual SQL Server machines to other ESX servers to allow the maintenance of ESX servers. Online patching for both vSphere hosts and Windows operating system is available using VMware Update Manager.

Understanding the patching and rollout in test and development environments as well as production environments will help you define how you should create, stage, and provision virtual machines in your data center, leveraging VMware Converter or the templates and cloning features in VMware vSphere.

Defining Capacity and System Growth Requirements

When designing your SQL Server consolidation project, ensure that you include headroom for future growth. Earlier we wrote about capturing SQL Server resource needs in terms of the usual system utilization metrics such as CPU, memory, and disk and network I/O, as well as user connections, transaction throughput, query execution efficiency and latencies, and database size. In addition, you might have established your own capacity management targets for system utilization on your SQL Server hosts, such as 70 percent CPU utilization, to leave enough headroom for any usage spikes. You must take this sizing and capacity planning into account along with plans for future growth when you design for deployments on VMware vSphere.

The good thing about virtualization is that VMware vSphere allows you to create a dynamically scalable SQL Server environment. Adding CPU, memory, and network resources to virtual machines can be done by simply clicking an Add hardware button and rebooting the virtual machine (Figure 7.10). Storage can be added without a virtual machine reboot. With vSphere, depending upon guest operating support, you can even hot add CPU or memory to scale up SQL Server virtual machines as needed, without the need to reboot.

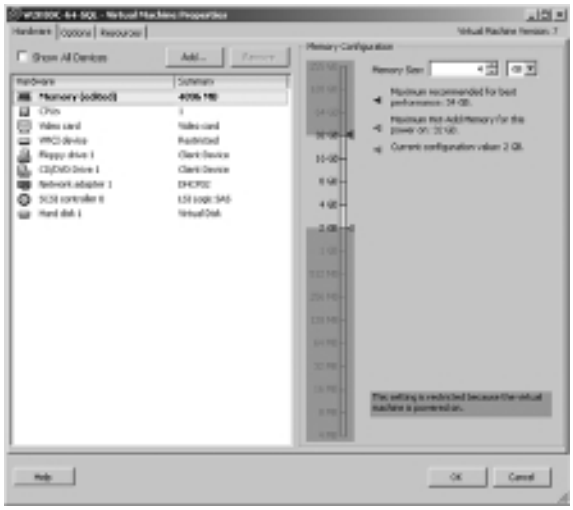

FIGURE 7.10 Clicking Add button on Hardware tab

Ranking SQL Server Candidates for Virtualization

Because SQL Server supports a very wide variety of applications, it is important to understand the business and technical requirements for SQL Server databases that support these applications. These databases typically differ based on various factors such as mission criticality, user count, database size, resource utilization, response time requirements, and workload profile. As a result, they often have unique deployment needs. You may have used these requirements to define service-level agreements (SLAs) and classify databases into multiple tiers. For example:

- Mission-critical databases, or so-called tier-1 databases that often have SLAs, require very high availability and minimal yearly downtime.
- Resource-intensive databases used by large number of users often need dedicated hardware for deployment.
- Other databases busy only during specific periods for such tasks as reporting, batch jobs, or ETL workloads may have much less stringent requirements for availability, scalability, number of users, and other similar factors.
- Many small lightly used databases that usually support business unit level or departmental applications have no high-availability configuration.

Understanding such classification can help you rank individual SQL Server databases for virtualization. You can prioritize the databases to virtualize starting with smaller, less risky ones first, and working your way to the more resource-intensive, highly available mission-critical databases.

Because you consolidate multiple workloads on a single ESX host, this ranking also helps you design in a vSphere environment in a couple of ways:

- Design an ESX and storage hardware configuration that meets all the resource requirements of the multiple workloads deployed together.
- Identify how you can leverage advanced VMware vSphere features such as DRS to optimize overall resource utilization and virtual machine performance.

Defining Deployment Strategy

After you establish baseline profiles for your existing SQL Server databases and shortlist the SQL Server databases you want to virtualize, the next step is to design a VMware vSphere architecture that satisfies these profiles. A key discussion at this time is to decide on a SQL Server virtual machine placement strategy.

Scale-up or Scale-out Strategy

A *scale-out* approach with multiple small SQL Server virtual machines (for example, one SQL Server database per virtual machine) is generally more appropriate for VMware vSphere. The scale-out approach makes it easier for you to customize the configuration as needed. As a result, this approach usually provides better workload and security isolation, easier maintenance, and more granular change management.

VMware DRS functions more effectively with smaller virtual machines, with the added benefit of faster VMotion migrations. Smaller virtual machines can be easily migrated by DRS in order to balance workload across cluster. Virtual machines with more memory or more vCPUs add

more constraints on placing virtual machines. These DRS considerations add one more reason to configure virtual machines with only as many vCPUs and only as much virtual memory as they need.

A *scale-up* approach, which involves a large virtual machine with multiple databases per SQL Server instance, usually exhibits higher-performance overheads associated with the scaling of SMP virtual processors. The scale-up approach can also result in bottlenecks in the operating system itself, especially if you run into the operating system's 32-bit memory limits. Also, as we stated earlier, moving larger virtual machines is more challenging for VMware DRS because it places greater demands on available resources.

Dedicated ESX Server for SQL Server Virtual Machines or Mixed Workload

Another commonly asked question is whether one should dedicate all SQL Server virtual machines to a single or a group of ESX servers. This question usually stems from the fact that database servers are typically considered more critical than other application servers. The short answer to this question is that you should enable VMware DRS and let it optimally balance the virtual machines running on the ESX servers. The continuous real-time monitoring performed by DRS not only helps in the initial placement of virtual machines but also performs the necessary VMotion migrations to keep resource utilization balanced across ESX cluster. We have talked to several customers who are successfully using DRS even with SQL Server virtual machines.

If you do prefer to place SQL Server virtual machines manually, the answer to this question will depend upon on both technical and operational considerations. SQL Server performance issues usually result from disk I/O bottleneck. Consolidating multiple SQL Servers with similar database workloads on a single ESX server will only aggravate the disk I/O bottleneck. In general, placing multiple virtual machines that bottleneck on similar resources is not recommended. Try to place workloads with similar I/O patterns on the same ESX servers. This will make it easier to create I/O profiles for your databases. Storage plays a paramount part in SQL Server database performance, and consolidation of like I/O patterns will yield better performance from your storage infrastructure. For example, placing a virtual machine performing sequential read I/O with another one doing random read I/O will result in random read I/O at the storage level and will negate any readahead benefits you expected to achieve.

Mixing different workloads that stress CPU, memory network, and disk I/O can optimize the ESX server hardware resources relatively better. One common benefit that is often cited is in-memory network communication between virtual machines colocated on the same ESX server. If the components of the application stack that are communicating with each other are placed on the same ESX server, the network traffic takes place in memory and does not have to go through physical network adapters. VMware studies have shown that this in-memory network traffic can exceed the typical 1 Gbps physical network bandwidth. Given these two factors, we suggest that you try to mix the SQL Server virtual machines strategically with other virtual machines.

Placing heterogeneous workloads into the SQL Server I/O path will require much more stringent planning and analysis for your overall workloads and system requirements. Workloads that have different I/O requirements from your SQL Server workload requirements might not be well suited for placement in a SQL Server consolidation environment, especially if your ability to test/model such workloads becomes challenging. In any case, best practices dictate that you separate non-SQL applications onto separate I/O paths.

Number of SQL Server Virtual Machines per vSphere Host

Customers often ask us how many SQL Server virtual machines they should run on a single vSphere host. The answer to this will primarily depend upon the actual resources your SQL Server databases use. This number will also be affected by other such requirements as capacity headroom, availability considerations, and so on. However, the key factor is the actual resource usage, such as CPU or memory utilization or disk I/O throughput, not simply the capacity of the physical hardware deployed for the physical SQL Server.

You can place multiple SQL Server virtual machines on a single vSphere host, but only a few of these may need to be active at a time. For example, the OLTP workloads may be dominant during business hours, and batch or backup jobs may be consuming resources during off-hours. If you can place your SQL Server VMs smartly, you can use VMware DRS features to favor OLTP or DSS workloads as needed. During the batch or backup windows, you can set high shares for a specific database and low shares for OLTP workloads likely to be least used at that time. So long as the physical hardware and shared storage can support the actual resources used by the SQL databases, you can run multiple SQL Server VMs per vSphere host. If you over-commit the actual resources being used beyond these physical resources, you will notice performance degradation.

Migrate to Virtual Environment

Several options are available to move your databases to the virtual environment. The good thing is that most of these options are the same ones you would use to move your databases in a physical environment.

You can use options specific to virtual environments such as the following:

- VMware Converter
- Templates and clones

You can also use the SQL Server–provided options, such as the following:

- Backup and restore
- Detach and attach
- Log shipping
- Database mirroring

VMware Converter

When you create the SQL Server virtual machine, you can simplify the physical to virtual migration using VMware Converter, which can both hot- and cold-clone an existing physical SQL Server instance (Figure 7.11). You should analyze and consider the validity of your physical server before using it as a candidate for conversion using VMware Converter. Remember that everything that currently exists on the physical server will be converted into the VM instance. Physical servers often have services and processes, such as SNMP monitoring and backup agents, that may be running in the background. When creating virtual machines using VMware Converter, be careful about such processes. These processes can add significant virtualization overhead and seriously

degrade SQL Server performance in a virtual environment. Monitor any extra processes running in a virtual machine and stop them.

FIGURE 7.11
VMware Converter

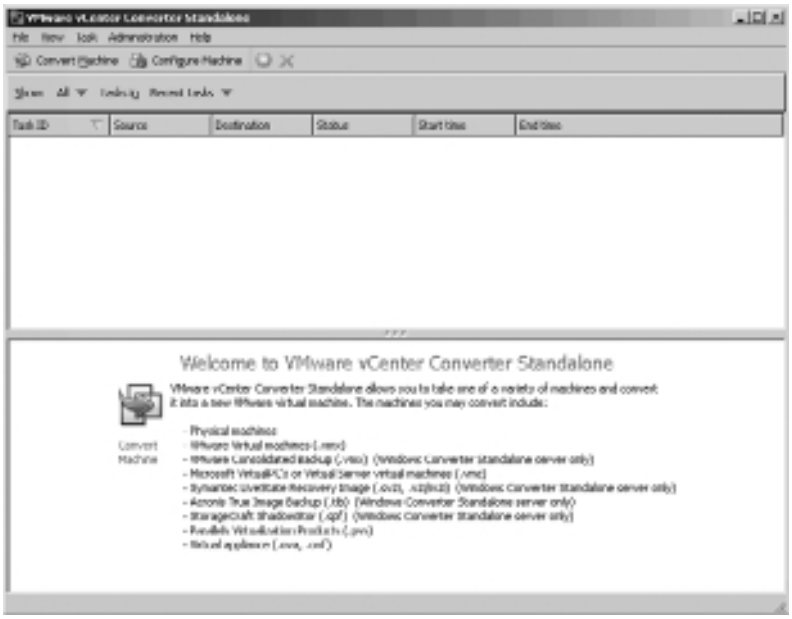

VM TEMPLATES

When creating SQL Server virtual machines, take advantage of templates and cloning features in VMware vSphere. We discussed this in detail in Chapter 3. Using templates, you can build a set of common server images with the required Windows Server OS patches and updates. You can then customize these templates further by scripting the SQL Server installation and locking down the SQL Server configuration as needed. You can easily clone or deploy these templates over and over again to reduce the time to provision SQL Server virtual machines. Creating such a clean installation of SQL Server, as compared to a physical to virtual conversion, ensures that there are no driver, security, and legacy software issues that might affect the SQL Server virtual machine.

To ease the initial transition to a virtual environment, you can configure SQL Server virtual machines with RDM in physical mode that point to existing data, transaction logs, and TempDB LUNs. Starting with RDM means one less variable to worry about during the migration project, but we do recommend a full transition to VMFS partitions for maximum flexibility.

BACKUP AND RESTORE

Backup and restore is the common method used to move a database from a physical environment to a virtual machine. This method requires that you take a full database backup and copy it to the virtual machine where it is restored. To make sure that you have completely synchronized database in the virtual machine, you will need to stop all active connections on the source database. Backup and restore works well for smaller databases but can be the most

time-consuming method for larger databases. Also, you will need to stagger your backup and restore tasks so that the I/O subsystem can handle the entire load.

Detach and Attach

Many DBAs choose to simply detach the database from the source server and reattach it onto the virtual SQL Server. If you are using SAN-based LUNs, this can be the fastest method for migrating databases to a virtual machine. As with backup and restore, the time to move the database files will increase with larger databases. This method also needs downtime to detach a database until the database files are copied to the virtual machine.

Log Shipping

You can also leverage SQL Server high-availability features such as log shipping and database mirroring for moving databases to virtual environment. These two methods require much smaller downtime compared to the backup and restore and detach and attach methods. In log shipping, you take a point-in-time full backup of the source database and then restore it to a destination server, the virtual machine in this case, using the `WITH STANDBY` or `WITH NORECOVERY` option. If you have log shipping enabled, all you need to do is make sure that the transaction logs have been shipped and applied in the virtual SQL database. With log shipping properly in place, you will need only to switch to the virtual SQL deployment. In addition, log shipping is relative simple to set up compared to database mirroring.

Database Mirroring

Like log shipping, database mirroring can migrate databases to virtual environment with minimal downtime. Database mirroring allows you to set up a mirror of a database on a different server, the virtual machine in this case. Once you have mirroring configured, the mirroring mechanism will keep the database in the virtual machine current with the source server. To transition to the virtual environment, it is a simple matter of failing over the mirrored environment. This failover will configure the mirror server as the new principal, reconnect many applications automatically to the new principal, and make the current principal into the mirror. You can continue to keep the database on the old server until you are satisfied with the virtual SQL Server. When you are comfortable with the virtual deployment, you can simply break the mirroring and remove the old database from the old server. Configuring database mirroring is relatively more complex compared to log shipping and has strict compatibility requirements across SQL Server versions.

Virtualizing SQL Server Proof of Concept

Source Discovery has decided to extend the success of its virtualization projects to other server applications. With the successful virtualization of Windows Servers, Active Directory, and Internet Information Services, Source Discovery is confident it can now tackle SQL Server virtualization. It has captured the physical resource utilization and SQL Server–specific performance metrics using the VMware Capacity Planner tool. It has prioritized its tier-2 and departmental SQL Server instances for virtualization in the first phase and plans to work its way to mission-critical tier-1 databases in the next phase. Most of these stand-alone SQL Server instances are

less resource intensive, host smaller databases with fewer users, and have no high-availability configuration. Source Discovery hopes that it can leverage VM templates to minimize the administrative and maintenance efforts. Also, it wants to leverage VMware DRS and HA to extend the uptime for these databases.

Source Discovery is also investigating the possibility of standardizing on Microsoft Office SharePoint Server as a collaboration suite. This SharePoint Server farm will be backed by a SQL Server configuration that needs to be highly available using failover clustering. Source Discovery has decided to adopt a "virtualize first" policy for this SharePoint Server deployment, so we need to build a clustered SQL Server configuration for the SharePoint Server farm.

In this example, you will start with the Windows Server 2008 templates from Chapter 3 to build the stand-alone as well as clustered SQL Server virtual machines. As you deploy these servers, you will configure the vCPU, memory, disks, and networking appropriately to match the applications requirements. We assume you have a thorough knowledge of SQL Server 2008, so we will not show all aspects of SQL Server configuration in this proof of concept. We will start with the stand-alone SQL Server VM configuration and then move to the clustered SQL Server configuration. We will rely on the details in Chapter 3 to build the basic failover cluster and provide the SQL Server–specific steps here.

Configuring vSphere

Before we get started with SQL Server deployment, follow this checklist to make sure ESX server best practices are configured:

- Make sure that the ESX server network interfaces for the Service Console, VMotion, and VMnetwork are teamed and are plugged into a fully redundant physical network infrastructure for no single point of failure at the network layer.
- Make sure VMware DRS and HA are enabled for the ESX cluster so that SQL Server virtual machines can benefit from higher availability (Figure 7.12).

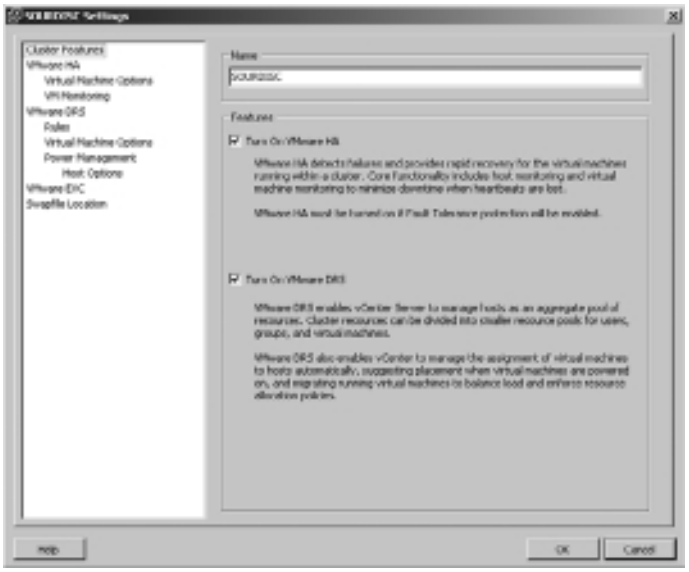

FIGURE 7.12
Enabling DRS and HA settings for vSphere cluster

- Increase the transmit coalescing size in the virtual NIC to support high network throughput demands. To increase transmit coalescing size, perform the following steps:
  1. In vSphere Client, choose the ESX host where the virtual machine is deployed.
  2. Select the Configuration tab.
  3. Click Advanced Settings in the Software panel.
  4. Click the Net tab.
  5. Edit the value of Net.vmxnetThroughputWeight, changing it from the default of 0 to 128, and then click OK (see Figure 7.13).
  6. Reboot the virtual machine.

FIGURE 7.13
Setting transmit coalescing for the vSphere host

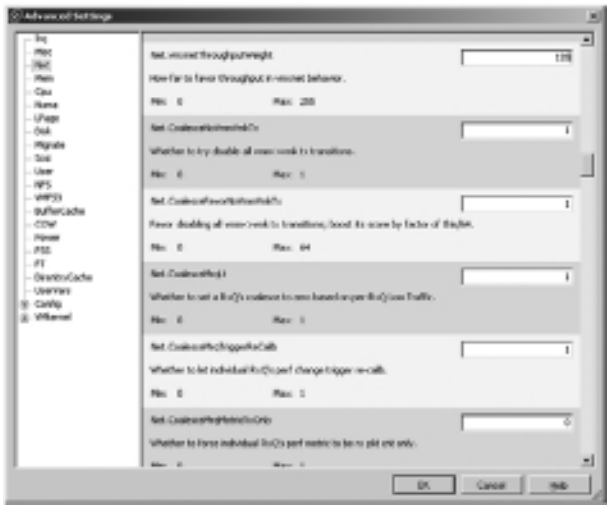

- Increase the number of outstanding disk I/O requests on the vSphere host to avoid any I/O command queuing. When virtual machines share a LUN, the total number of outstanding commands permitted from all virtual machines to that LUN is governed by the Disk.SchedNumReqOutstanding configuration parameter. Note that the setting has no effect when there is only one virtual machine issuing I/O to the LUN. To increase the number of outstanding disk I/O requests, perform the following steps:
  1. In vSphere Client, choose the ESX host where the virtual machine is deployed.
  2. Select the Configuration tab.
  3. Click Advanced Settings in the Software panel.
  4. Click the Disk tab.
  5. Edit the value of Disk.SchedNumReqOutstanding, changing it from the default of 32 to 64, and then click OK (see Figure 7.14).
  6. Reboot the virtual machine.

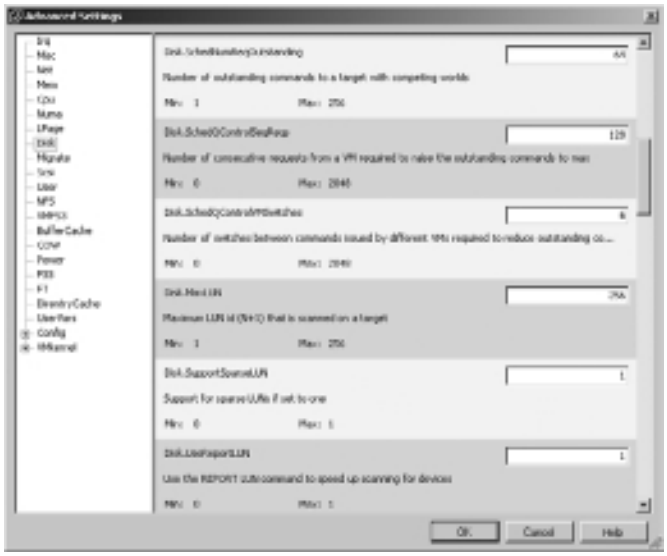

FIGURE 7.14
Setting the outstanding disk I/O requests for the vSphere host

Configuring SQL Server Virtual Machines

Based on the observed SQL Server workload in the physical environment, we plan to deploy a SQL Server 2008 virtual machine with following configuration:

- Two virtual CPU with 4 GB of memory
- One VMXNET3 network adapter
- Four virtual disks (OS, database, logs, and TempDB)
- Two virtual SCSI controllers (one for OS and one for the database, TempDB, and logs)

Figure 7.15 shows these settings for the SQL Server virtual machine.

As stated earlier, storage configuration is the key to any SQL Server deployment. We have configured the SQL Server virtual machine used for the proof of concept with four separate virtual hard disks. We have made sure that these virtual disks span LUNs that are carved out of separate physical disk drives in the storage array. Table 7.3 shows these details.

TABLE 7.3: SQL Server VM Virtual Disk Configuration Details

VM HARD DISK	SIZE	PURPOSE	VSCSI CONTROLLER
Hard disk 1	40 GB	OS boot drive	0:0 (LSI Logic SAS)
Hard disk 2	40 GB	SQL database LUN	1:0 (Paravirtual)
Hard disk 3	20 GB	SQL transaction log LUN	1:1 (Paravirtual)
Hard disk 4	20 GB	TempDB database LUN	1:2 (Paravirtual)

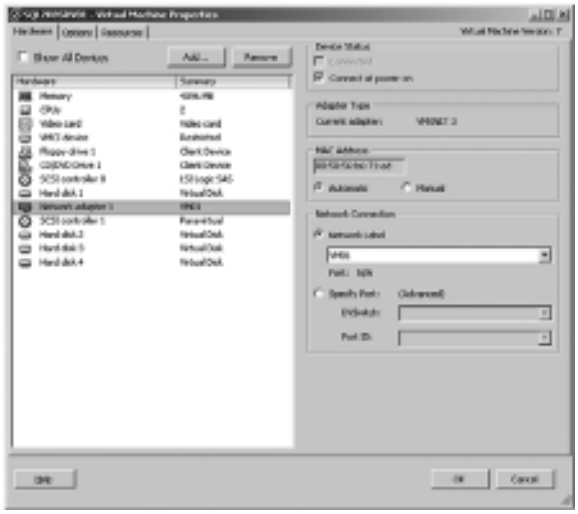

FIGURE 7.15
SQL Server VM, NIC configuration

At the time of this writing, vSphere 4 does not support PVSCSI adapters for boot partition. You have to create a primary adapter for use with a disk that hosts the system software (boot disk) and a separate PVSCSI adapter for the disk that stores user data, such as a database. For this deployment, we'll stick with the default LSI Logic SAS adapter for virtual machines with the Microsoft Windows 2008 guest operating system. PVSCI adapters are more efficient and recommended for I/O-intensive virtual disks.

Figure 7.16 illustrates the configuration for the PVSCSI controller.

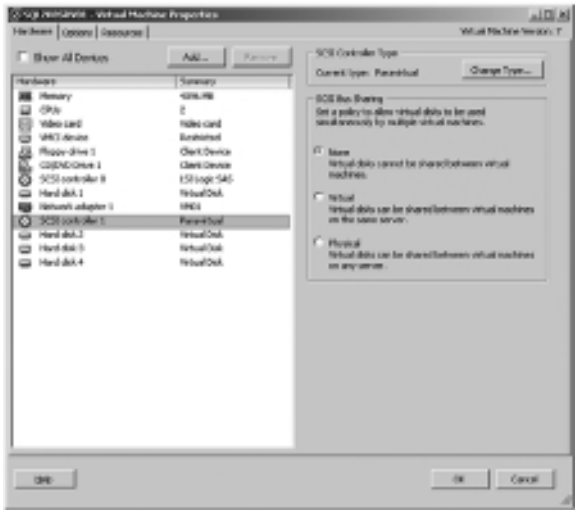

FIGURE 7.16
SQL Server VM, PVSCSI controller configuration

Paravirtual SCSI Controller as Boot Device

vSphere 4 Update 1 introduced the ability to use the VMware paravirtualized SCSI controller as the boot controller. Prior to this, using the PVSCSI controller as the boot device was not supported.

Figures 7.17 through 7.19 show the individual virtual hard disks to be used by the SQL Server database. All these virtual disks use thick provisioning to avoid any first-time allocation performance hit. Figure 7.17 shows the virtual hard disk to be used for data files.

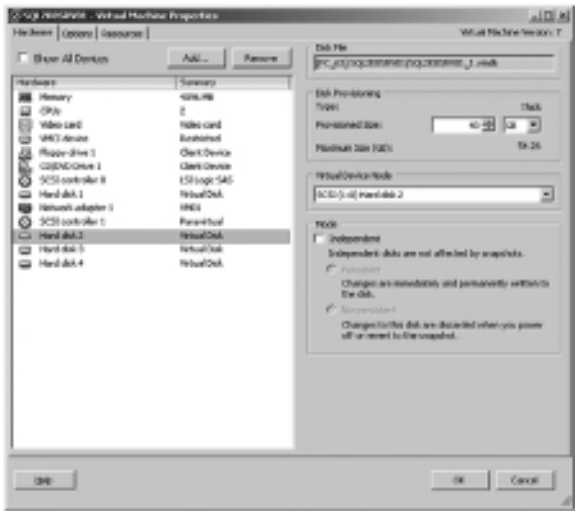

Figure 7.17
SQL Server VM, data LUN configuration

Figure 7.18 shows the virtual hard disk that will be used for transaction logs.

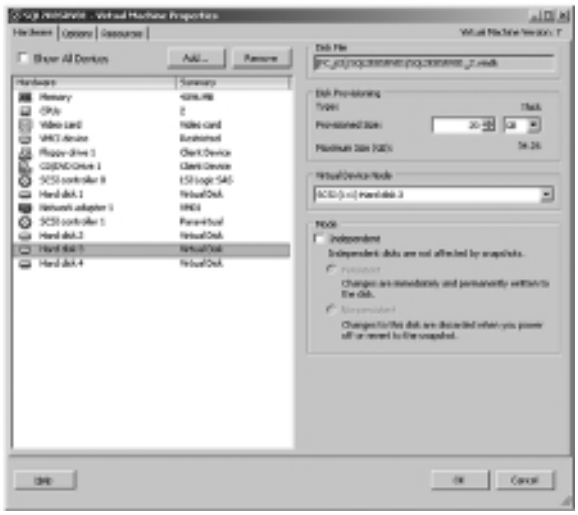

Figure 7.18
SQL Server VM, log LUN configuration

Figure 7.19 shows the virtual hard disk that will be used for TempDB.

FIGURE 7.19
SQL Server VM, TempDB LUN configuration

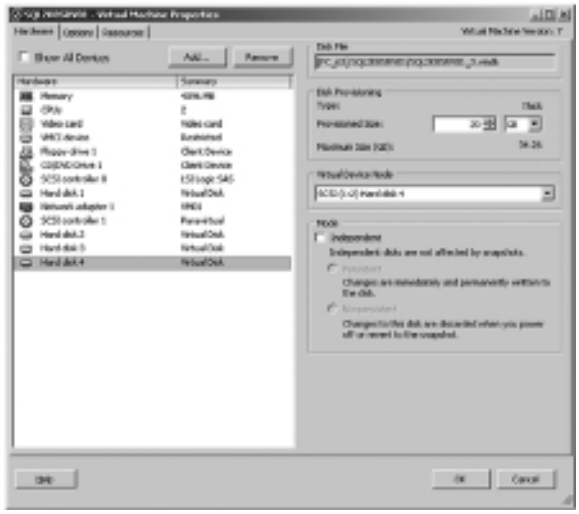

Aligning file system partitions is a well-known storage best practice for database workloads. Partition alignment for virtual machine disks and VMware VMFS partitions prevents performance I/O degradation caused by I/O crossing track boundaries. VMFS partitions created using vSphere Client are automatically aligned on 64 KB boundaries as recommended by storage and operating system vendors. VMware tests have reported that aligning VMFS partitions to 64 KB track boundaries results in reduced latency and increased throughput. If you have manually created the virtual disks using the vmkfstools utility, you need to make sure the partitions are aligned correctly. This will reduce the chance of split writes and avoid any performance degradation.

Partition Alignment Inside a Windows VM

Windows Server 2008 automatically aligns partitions using a default starting offset of 1 MB (1,048,576 bytes) and in many cases is the recommended setting; however, consult your storage vendor or SAN administrator to make sure these settings are suitable for your environment. In versions of Windows prior to Windows 2008, you need to align the partitions manually at the time of creation using `diskpar.exe` or `diskpart.exe`. The following illustration shows how you can use the DiskPart tool to format your disks with a 64-bit alignment for Windows Server 2003.

C:\>diskpart
Microsoft DiskPart version 5.0.6495
Copyright (C) 1999-2007 Microsoft Corporation.
On computer: CLOUDDB01
DISKPART> select disk 4
Disk 4 is now the selected disk.
DISKPART> create partition primary align=64
DiskPart succeeded in creating the specified partition.
DISKPART> assign letter=L
DiskPart successfully assigned the drive letter or mount point.
DISKPART> _

The Disk Management console display should be similar to Figure 7.20 after all of your disks have been properly formatted. After formatting is complete, your disks are ready for use. You can now close the Server Manager utility.

FIGURE 7.20
Disk Manager display

You can now proceed with SQL Server installation (Figure 7.21). The SQL Server installation process in the virtual machine is similar to that in a physical environment. If needed, it will install an updated version of the .NET Framework and Windows Installer as prerequisites. You will install the SQL Server software in the suggested folders on the C: drive.

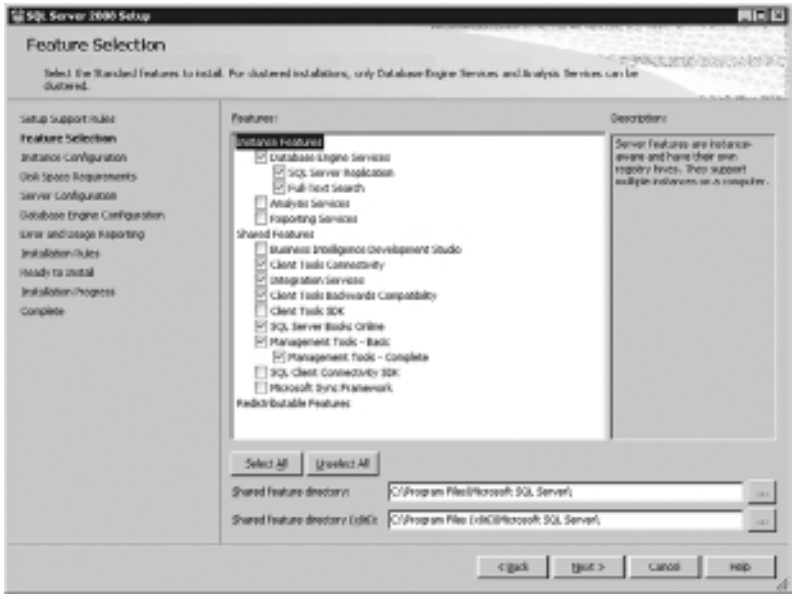

FIGURE 7.21
SQL Server installation

For the data directories, make sure you point to the virtual disks that you added to the virtual machine (Figure 7.22).

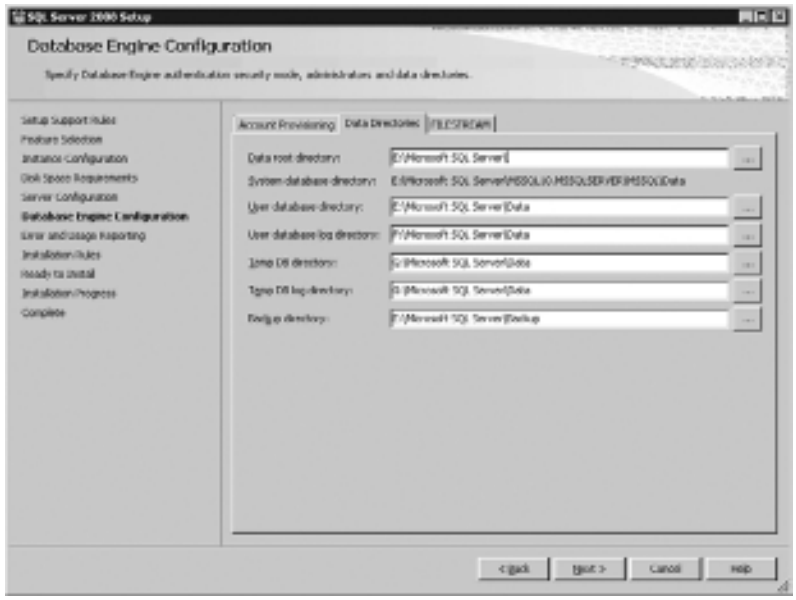

FIGURE 7.22
SQL Server installation, data directories' location

The rest of the SQL Server installation is the same as in the physical environment. Figure 7.23 shows the successful completion of the SQL Server installation.

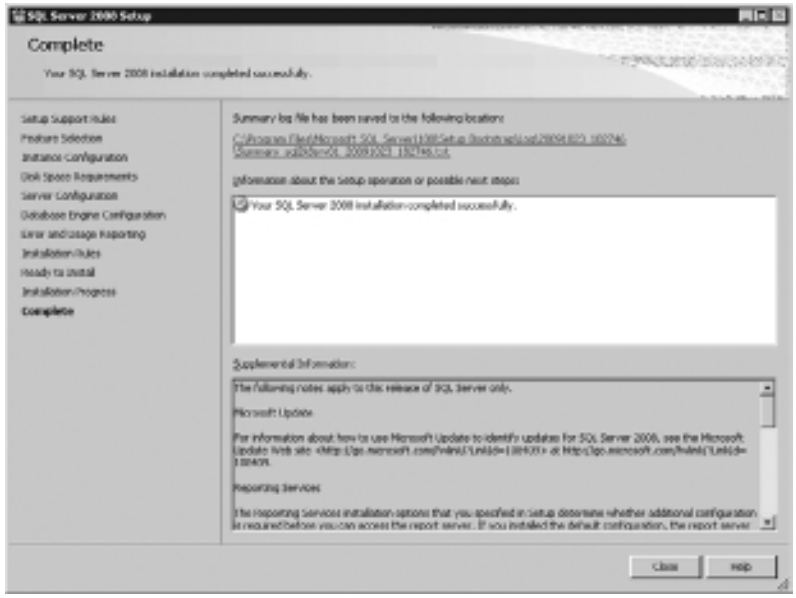

FIGURE 7.23
Successful SQL Server installation

Next we will show how to configure some SQL Server–specific settings to optimize its performance. In particular, we will configure the following:

- Administrator account that starts the SQL Server service with the Lock Pages In Memory right
- SQL Server to use large memory pages

Locking SQL Server Buffer Pool Pages in Memory

Caching data in the SQL Server buffer cache to avoid the disk I/O performance penalty is a common optimization technique for SQL Server. You can optimize this further by preventing the Windows operating system from paging out the buffer pool memory to the page file on disk. You achieve this by locking the memory that is allocated for the buffer pool in physical memory. You lock the memory by assigning the Lock Pages In Memory right to the user account that is used as the start-up account of the SQL Server service.

To assign the Lock Pages In Memory right, follow these steps (Figure 7.24):

1. Click Start, click Run, type **`gpedit.msc`**, and then click OK. The Group Policy dialog box appears.
2. Expand Computer Configuration, and then expand Windows Settings.
3. Expand Security Settings, and then expand Local Policies.
4. Click User Rights Assignment, and then double-click Lock Pages In Memory.
5. In the Local Security Policy Setting dialog box, click Add User Or Group.
6. In the Select Users Or Groups dialog box, add the account that has permission to run the `Sqlservr.exe` file, and then click OK.
7. Close the Group Policy dialog box.
8. Restart the SQL Server service.

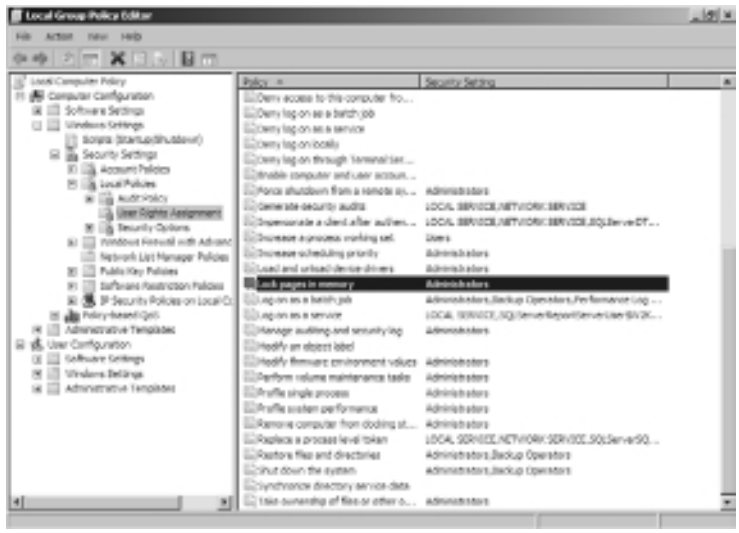

FIGURE 7.24
Lock Pages In Memory right

NOTE For 64-bit editions of SQL Server, initially only SQL Server Enterprise edition could use the Lock Pages In Memory right. For SQL Server Standard editions, the support for the Lock Pages In Memory right has been introduced in later Service Pack cumulative updates. For more information on this support, view article 970070 in the Microsoft Knowledge Base.

After you assign the Lock Pages In Memory right, restart the SQL Server service. You can validate that the user right is used by the instance of SQL Server by making sure that the following message is written in the SQL Server Error Log at startup: "Using locked pages for buffer pool."

When you enable the Lock Pages In Memory right, you should also configure the Memory Reservation setting for the SQL Server VM using Virtual Machine Properties window. This reservation setting guarantees the amount of memory for the VM and will ensure that the VM memory is not ballooned or swapped by vSphere.

Using Large Memory Pages with SQL Server

SQL Server also provides various trace flags to improve performance when running in high transaction throughput workloads. One of the trace flags, trace flag 834, causes SQL Server to use Microsoft Windows large-page allocations for the memory that is allocated for the buffer pool. Normally Windows allocates memory in 4 KB pages. When large pages are enabled, the page size may be from 2 MB to 16 MB depending on the hardware platform. Large page support can be enabled for 64-bit SQL Server Enterprise edition when the virtual machine memory is 8 GB or more and the Lock Pages In Memory privilege is set.

Large pages are particularly useful when virtualizing SQL Server using the newer processors that support hardware assist for MMU virtualization. These processors typically improve the performance for many workloads but can also introduce increased latency if the page table translations are not found in the CPU translation look-aside buffer (TLB). Using large pages in SQL Server backed up by large pages in ESX can eliminate or mitigate this cost of processing any TLB misses. Large pages can therefore reduce overhead because of the hardware MMU in addition to the performance gains they provide to many applications.

NOTE There are caveats to using large pages such as you can turn on trace flag 834 only at start-up. Large pages are allocated at start-up and are kept throughout the lifetime of the process. Trace flag 834 may significantly delay or prevent SQL Server from starting if memory is fragmented and if large pages cannot be allocated. Therefore, trace flag 834 is best suited for servers that are dedicated to SQL Server and should be enabled only with careful consideration of settings such as Max Server Memory. As with other advanced features, you should test this feature to validate whether it really benefits you.

Setting Up SQL Server Failover Clustering

In Chapter 3, we showed how to configure a Windows 2008 failover cluster across physical hosts using virtual machines. We assume that you have read all those steps, because we will leverage those steps to configure the SQL Server failover cluster in this section. As mentioned earlier, we plan to use this SQL Server failover cluster to host the Microsoft SharePoint proof of concept that we configure in the next chapter. We will show how to build a Microsoft Windows failover cluster that is configured as a clustered across box.

Before you configure the cluster, you must configure all the virtual disks you plan to use for the SQL Server VM. We plan to separate the different SharePoint Server databases, which is why we need so many virtual disks. Table 7.4 shows these details.

TABLE 7.4: Clustered SQL Server VM Virtual Disk Configuration Details

VM Hard Disk	Drive	Size	Purpose	vSCSI Controller
Hard disk 1	C:	32 GB	OS boot drive	0:0 (LSI Logic SAS)
Hard disk 2	E:	16 GB	Program files	0:1 (LSI Logic SAS)
Hard disk 3	Q:	1 GB	QUORUM LUN	1:0 (LSI Logic SAS)
Hard disk 4	M:	4 GB	DTC LUN	1:1 (LSI Logic SAS)
Hard disk 5	T:	40 GB	TempDB database LUN	1:2 (LSI Logic SAS)
Hard disk 6	L:	40 GB	TempDB logs LUN	2:0 (LSI Logic SAS)
Hard disk 7	D:	260 GB	ContentDB database LUN	1:3 (LSI Logic SAS)
Hard disk 8	F:	30 GB	ContentDB transaction log LUN	2:1 (LSI Logic SAS)
Hard disk 9	G:	200 GB	SearchDB database LUN	1:4 (LSI Logic SAS)
Hard disk 10	H:	30 GB	SearchDB transaction log LUN	2:2 (LSI Logic SAS)
Hard disk 11	I:	40 GB	ConfigDB database LUN	1:5 (LSI Logic SAS)
Hard disk 12	J:	20 GB	ConfigDB transaction log LUN	2:3 (LSI Logic SAS)

As discussed in Chapter 3, we have chosen an LSI Logic SAS adapter and changed SCSI bus sharing to physical mode for the virtual disks that will be part of the SQL Server failover cluster. We added three vSCSI disk controllers since separate disk controllers are required for the operating system disks and the clustered disks. Here we have employed a third vSCSI disk controller to separate the database disks from the transaction log disks. This will provide the ability to separate the database I/O from the transaction log I/O. Databases use random I/O patterns, whereas transaction log I/O uses sequential I/O patterns. The separation of these I/O patterns result in more efficient utilization of the storage array.

Figure 7.25 shows the configuration of the virtual machine for the first node of the SQL Server cluster. In this illustration, you can also see that we have provided the virtual machine with four virtual CPUs (vCPUs), 8 GB of virtual machine memory, and two virtual network interface cards (vNICs) consisting of a public network interface for client access and cluster management and a heartbeat network interface for cluster traffic.

FIGURE 7.25
SQL cluster node 1 VM properties

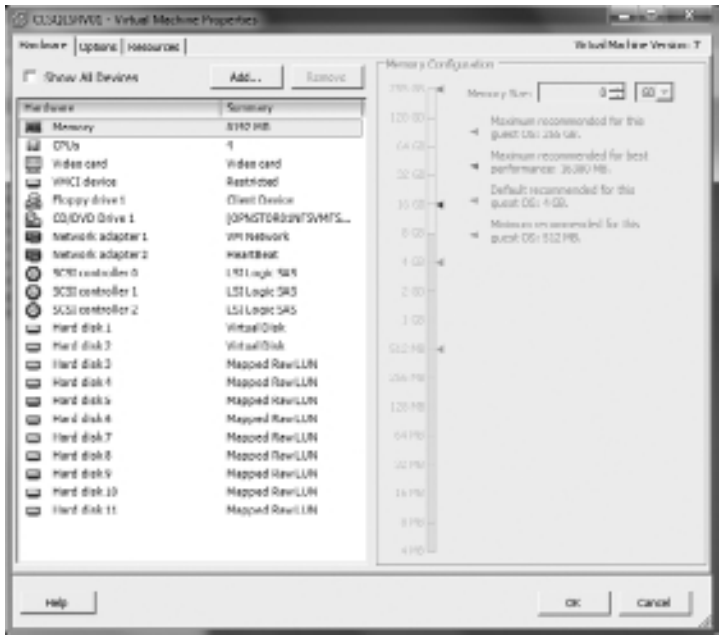

This configuration is duplicated for the second SQL Server node. Figure 7.26 shows the configuration of the virtual machine for the second node of the SQL Server cluster.

FIGURE 7.26
SQL cluster node 2 VM properties

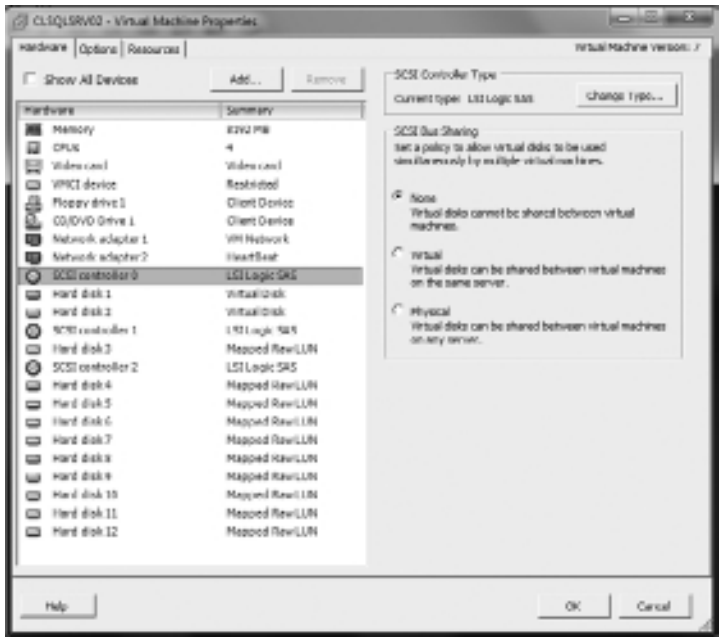

You need to boot node 1, mark all disks as online, and initialize and format them. You also need to repeat these steps on node 2. Figure 7.27 and Figure 7.28 show Server Manager on each node after these steps.

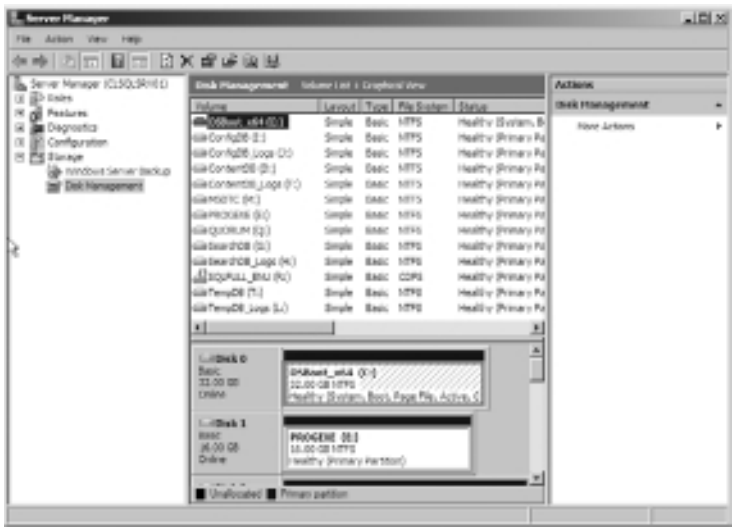

FIGURE 7.27
SQL cluster node 1
Server Manager

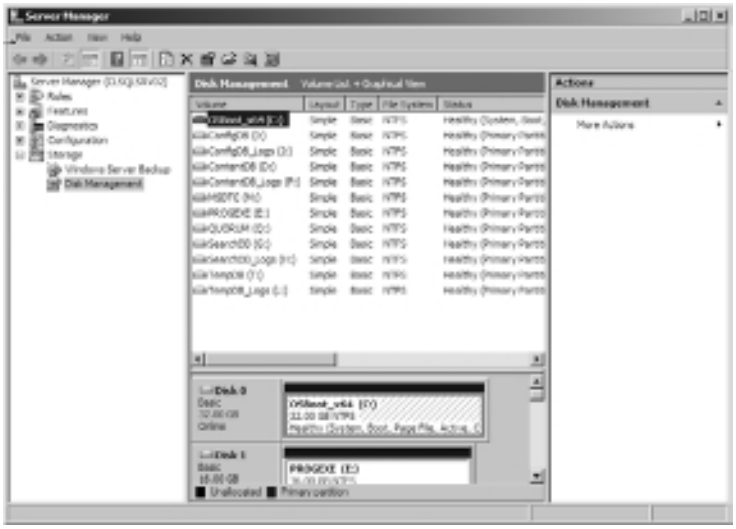

FIGURE 7.28
SQL cluster node 2
Server Manager

Before you proceed to create Windows failover clusters, change the VMware HA and DRS settings to ensure that SQL Server failover clusters function correctly in a vSphere cluster with HA and DRS enabled. You need to make the following changes for the failover cluster virtual machines:

1. Create DRS anti-affinity rules.
2. Enable strict enforcement of the anti-affinity rules you created in the previous step.

3. Disable the DRS automation level.
4. Disable the virtual machine restart priority.

DRS affinity rules can keep virtual machines on the same ESX host or make sure they are always on different hosts. For our proof of concept, we are deploying the CAB configuration, so you need to create an anti-affinity rule to make sure that the SQL Server nodes in the failover cluster do not run on a single ESX host. To create the anti-affinity rule, right-click the cluster name, and select Rules under VMware DRS in the Cluster Settings dialog box. Make sure you select Separate Virtual Machines as the rule type for the anti-affinity rule (see Figure 7.29).

FIGURE 7.29
Creating a DRS anti-affinity rule

To ensure that the anti-affinity rule you created is strictly applied, you need to configure an advanced option for VMware DRS. In the left pane of the Cluster Settings dialog box, select VMware DRS, and click Advanced Options. You need to manually enter **ForceAffinePoweron** in the Option column and **1** in the Value column (Figure 7.30).

FIGURE 7.30
Enabling strict enforcement of the anti-affinity rule

Virtual machines running failover clusters cannot be migrated using VMotion. So, you must disable the automation level of all virtual machines in the proof-of-concept SQL Server failover

cluster. In the Cluster Settings dialog box, select Virtual Machine Options in the left pane under VMware DRS. Make sure to select the Enable Individual Virtual Machine Automation Levels check box and then for each virtual machine in the SQL Server failover cluster, select Disabled from the drop-down menu in the Automation Level column (Figure 7.31).

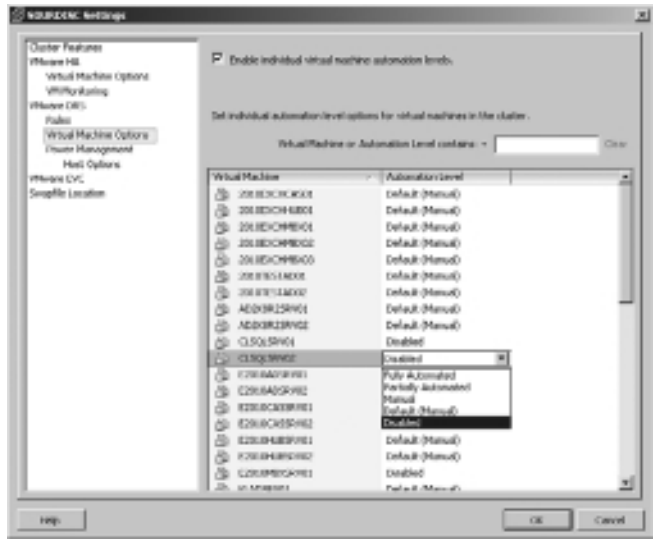

FIGURE 7.31
Disabling the virtual machine DRS automation level

VMware HA does not obey affinity or anti-affinity rules, so you must configure a different setting to address how the SQL Server failover cluster interacts with VMware HA. VMware HA provides a VM Restart Priority setting that determines the order in which virtual machines are restarted when the host fails. For failover cluster virtual machines, you need to disable the Restart Priority setting. In the Cluster Settings dialog box, select Virtual Machine Options in the left pane under VMware HA. For each virtual machine in the SQL Server failover cluster, select Disabled from the drop-down menu in the VM Restart Priority column (see Figure 7.32).

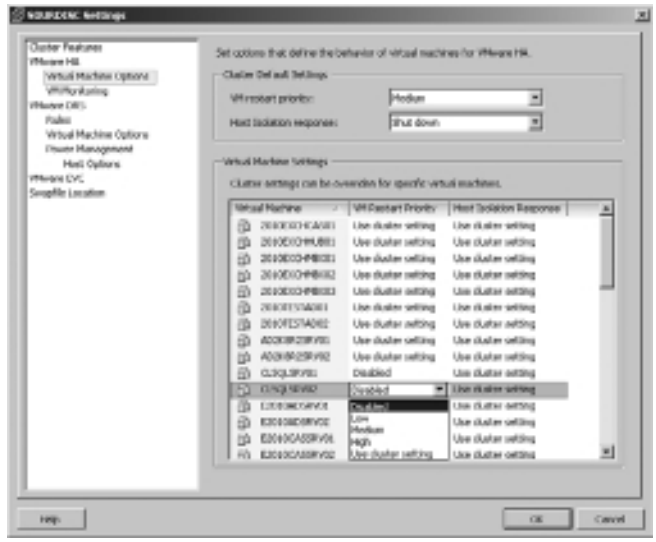

FIGURE 7.32
Disabling the virtual machine Restart Policy setting

Now you are ready to use the Windows failover clustering management utility to create the failover cluster. We detailed all the steps to do this in Chapter 3. Figure 7.33 shows the final cluster you built in preparation for the clustered SQL Server database back end.

FIGURE 7.33
Create Cluster Wizard

Make sure you configure the quorum virtual disk as the storage witness during cluster creation, as shown in Figure 7.34.

FIGURE 7.34
Cluster quorum settings

After the cluster creation is complete, validate the cluster to verify that all tests complete successfully, as shown in Figure 7.35.

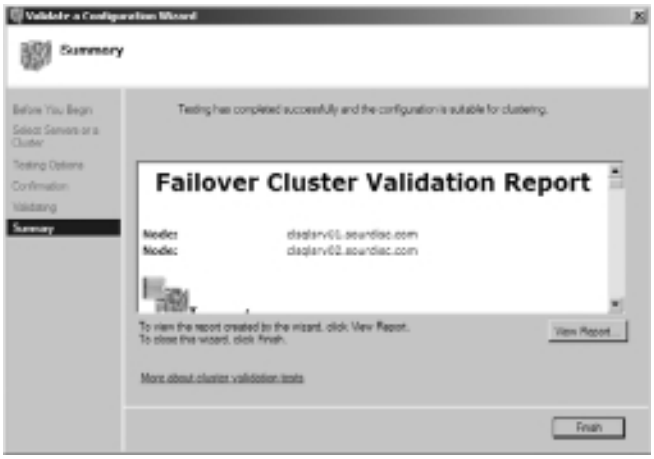

FIGURE 7.35
Cluster validation report

With the cluster configured and fully functional, you need to configure the Microsoft Distributed Transaction Coordinator (MSDTC). The SQL Server software will not install in a clustered configuration without the creation of the clustered DTC resource. Clustered SQL Server uses the MSDTC service for distributed queries, replication functionality, and two-phase commit transactions. Ensure that you have an available IP address for the MSDTC configuration.

To configure the MSDTC, open the Failover Cluster Management console, and right-click the cluster name. On the Server And Applications context menu, select Configure A Service Or Application. The High Availability Wizard Before You Begin screen is displayed. Click the Next button to continue with configuration. Select DTC from the list, and follow the steps to complete the configuration (Figure 7.36).

High Availability Wizard
Select Service or Application
Before You Begin
Select Service or Application
Client Access Point
Select Storage
Confirmation
Configure High Availability
Summary
Select the service or application that you want to configure for high availability:
DFS Namespace Server
DHCP Server
Distributed Transaction Coordinator (DTC)
File Server
Generic Application
Generic Script
Generic Service
Internet Storage Name Service (iSNS) Server
Message Queuing
Description:
DTC supports distributed applications that perform transactions. A transaction is a set of related tasks, such as updates to databases, that will either succeed or fail as a unit.
More about services and applications you can configure for high availability
< Previous
Next >
Cancel

FIGURE 7.36
Distributed Transaction Coordinator service setup

Figure 7.37 shows the SQLCluster.sourdisc.com Windows failover cluster we have built. You can also see the two nodes in this failover cluster as well as the ClusterDtc resource.

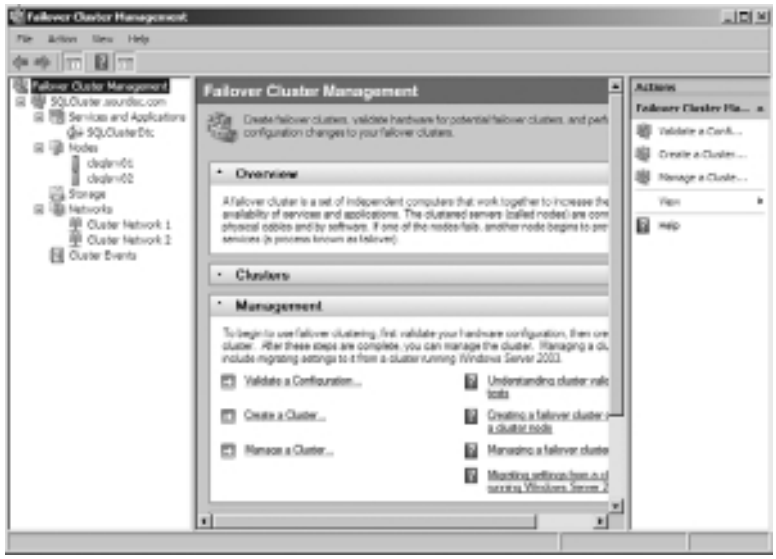

FIGURE 7.37
Windows failover cluster

Once you configure the Windows failover cluster with the clustered DTC resource, you install the SQL Server 2008 Enterprise edition in a clustered configuration. This is a two-part installation. You first install the software on the first node of the Windows failover cluster to create a single-node SQL failover cluster, after which you will repeat the install and add the second node of the SQL Server failover cluster. Figure 7.38 shows how you select the SQL Server failover cluster installation.

FIGURE 7.38
SQL Server Installation Center screen

The SQL Server failover installation will check all the necessary prerequisites to ensure that all the disks, cluster nodes, networks, and clustered DTC resources are in place before the actual software installation. You can see the clustered events log for any warnings and errors that might prevent the cluster from performing optimally. Figure 7.39 shows the set of support rules that are run.

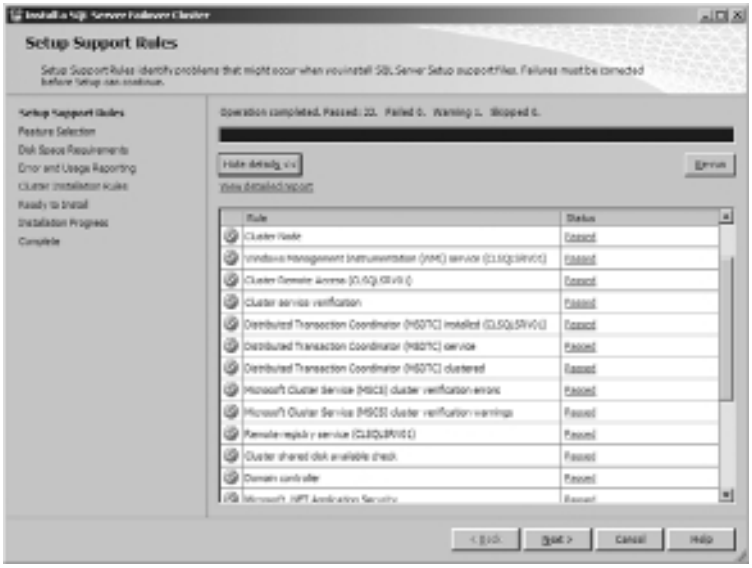

FIGURE 7.39
SQL cluster installation's Setup Support Rules page

Next you select the default cluster resource group. On the Cluster Disk Selection screen, configure all the database, logs, and backup directory locations, keeping in mind that all of these locations must be on the clustered shared disks (Figure 7.40). Make sure you select the correct disk resources that were created earlier.

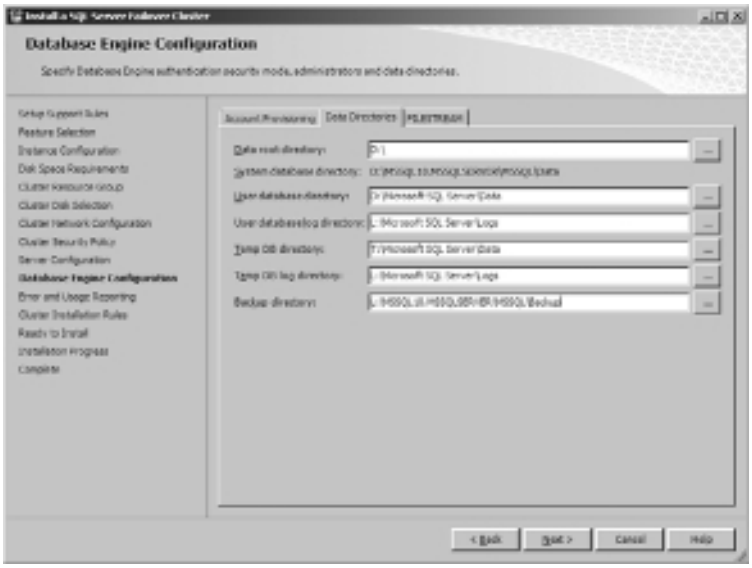

FIGURE 7.40
SQL cluster disk selection

For the cluster network configuration and security policy, select the default options. On the Cluster Installation Rules screen, you can see that all applicable tests passed without incident (Figure 7.41). If any test fails, view the failed information, correct the failure, and rerun the tests. Click the Next button to continue.

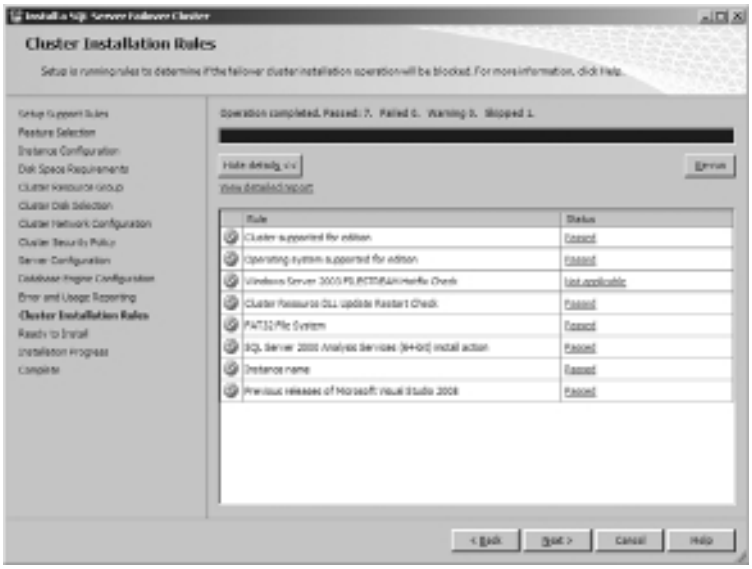

FIGURE 7.41
Cluster Installation Rules screen

Figure 7.42 shows the successful completion SQL Server cluster installation on the first node of the Windows failover cluster. Next you need to make that clustered SQL instance highly available by adding a second node of the Windows failover cluster to the SQL Server failover cluster configuration.

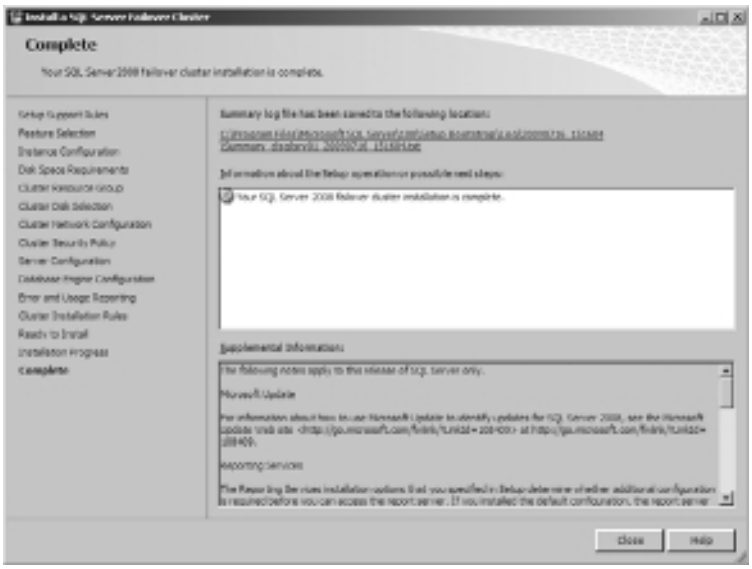

FIGURE 7.42
Successful completion of the first node

The procedure to add a second node to a SQL Server failover cluster is similar to the failover cluster install of SQL Server. You perform all of the same checks, and the installation will proceed as it did during the install for the first node of the SQL Server failover cluster. Figure 7.43 shows how you start the installation for the second node in the SQL Server failover cluster.

FIGURE 7.43
Adding the second node to the SQL Server failover cluster

Make sure all the setup support rules pass. The installer program should automatically pick the information that is relevant to the SQL Server failover clustered instance. The installation will also run a set of rules before adding the second node to make sure that all prerequisites are met (Figure 7.44).

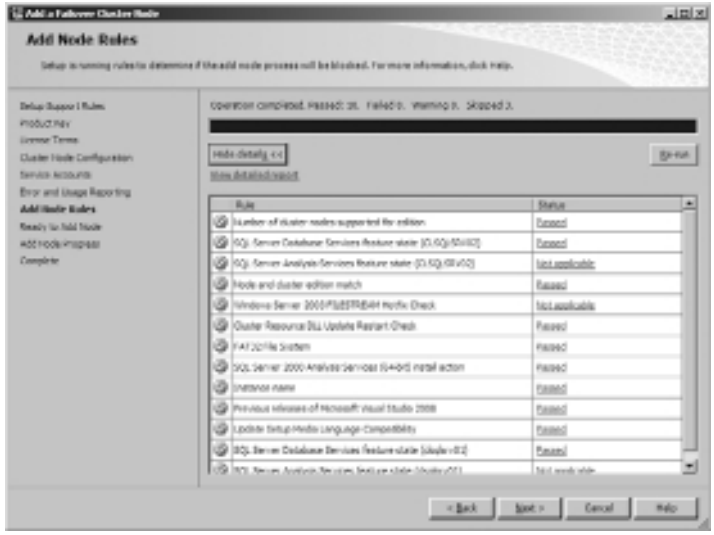

FIGURE 7.44
SQL Server second node rules

Just before the installation is completed, the SQL Server software checks that both nodes of the SQL Server failover cluster are running the same versions and patch levels. Once the installation has successfully completed, you have a two-node SQL Server failover cluster that you can use as the back-end database for the SharePoint Server farm. Figure 7.45 shows the successful completion of the second Windows failover cluster node to the SQL Server failover cluster.

FIGURE 7.45
Successful completion of the second node

With the SQL Server failover cluster successfully completed, you should perform instance failover tests to ensure that all resources and services failover properly.

Figure 7.46 shows that all resources and services are owned by cluster node CLSQLSRV01 and are online and functioning properly. Also, you can see that no warning events or errors are visible. Any errors or warnings that the cluster is experiencing will appear as red X or yellow exclamation point at the top of the cluster management utility. Next you perform a manual failover to the secondary node of the cluster CLSQLSRV02.

FIGURE 7.46
SQL Server failover cluster before failover

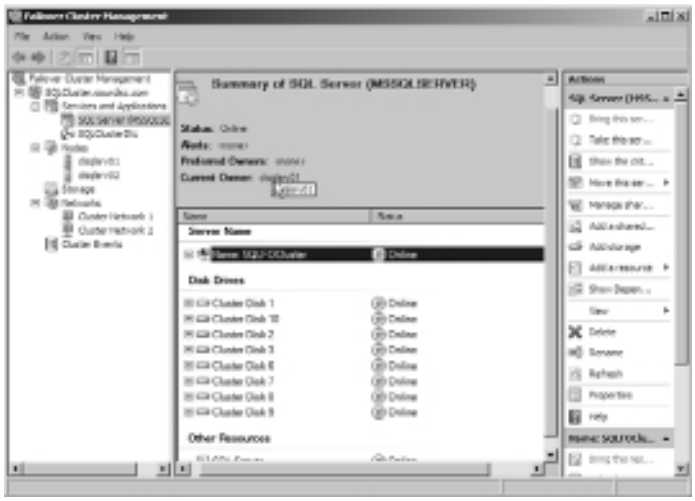

Figure 7.47 shows that the SQL Server services and its resources have been moved to the second node of the SQL Server failover cluster. You can see that all services and resources are online and functioning properly. In the next chapter, you will leverage this SQL Server failover cluster to configure the SharePoint Server farm.

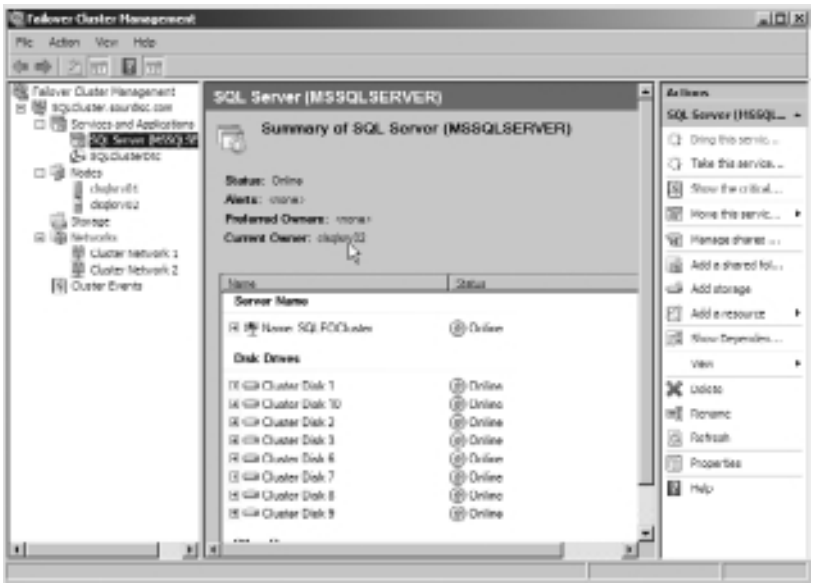

FIGURE 7.47
SQL Server failover cluster after failover

Testing Virtual SQL Server Deployment

Once your test environment is created, simulate all production environment activities, and monitor the environment. Test current production environment procedures and policies, and adjust and document changes where necessary to support the proposed consolidated environment. Add more workloads as appropriate, keeping in mind CPU, memory, network, storage (CORE 4), and other critical SQL Server metrics to ensure that the ESX server or its resources is not oversubscribed.

Several testing tools, free and commercial, are available to test your deployment. You can start with tools from Microsoft such as SQLIOSim or SQLIO. SQLIOSim is included with SQL Server 2008 and is located in the `MSSQL\BINN` folder (see Figure 7.48). You should run SQLIOSim on all I/O paths, ideally before you actually install SQL Server or right after to test for any data integrity issues in the underlying storage subsystem.

Update or add files to the SQLIOSim configuration to exercise the same I/O paths as the SQL Server database files. You should not point SQLIOSim to the same filenames as used by a SQL Server database because SQLIOSim can damage database files in a way that you cannot recover any data. We recommend you run the SQLIOSim test several times for reliable results (Figure 7.49). The SQLIOSim log file provides various details such as I/O throughput for each test that was run.

FIGURE 7.48
SQLIOSim tool

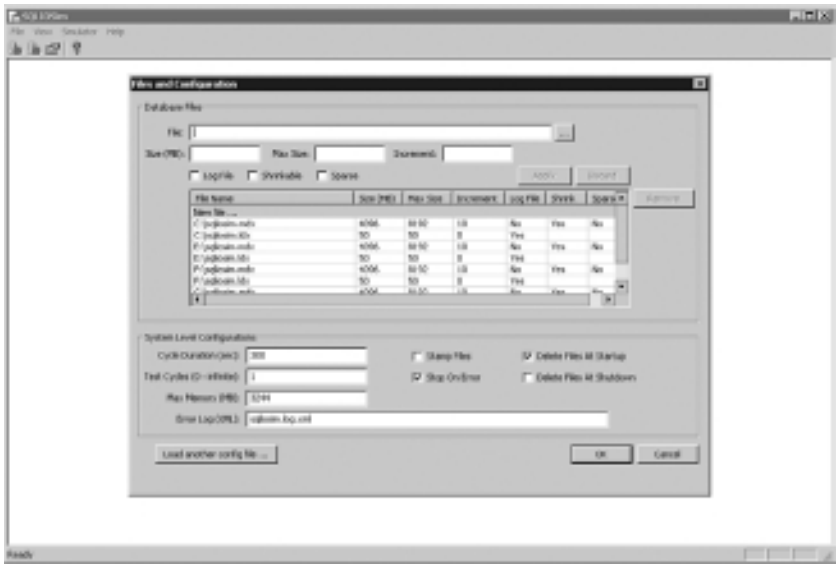

FIGURE 7.49
Running SQLIOSim test

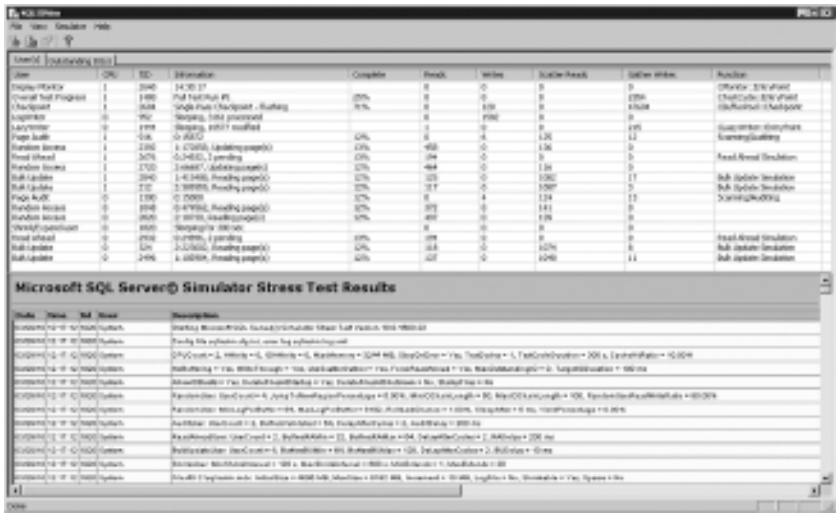

Note that Microsoft recommends using SQLIOSim only to check the stability of the I/O subsystem, and not as an I/O performance benchmarking tool. This is because SQLIOSim combines random I/O patterns together with known SQL Server I/O patterns to test the data integrity capabilities of your I/O path.

To test the throughput of the I/O susbsystem configured for SQL Server, you can use SQLIO or IOMeter. Both SQLIO and IOMeter can simulate various I/O patterns that can be used to match SQL Server data and log I/O behavior. VMware has frequently used IOMeter, shown in

Figure 7.50, for its storage performance studies and has published several white papers based on this tool.

FIGURE 7.50
IOMeter tool

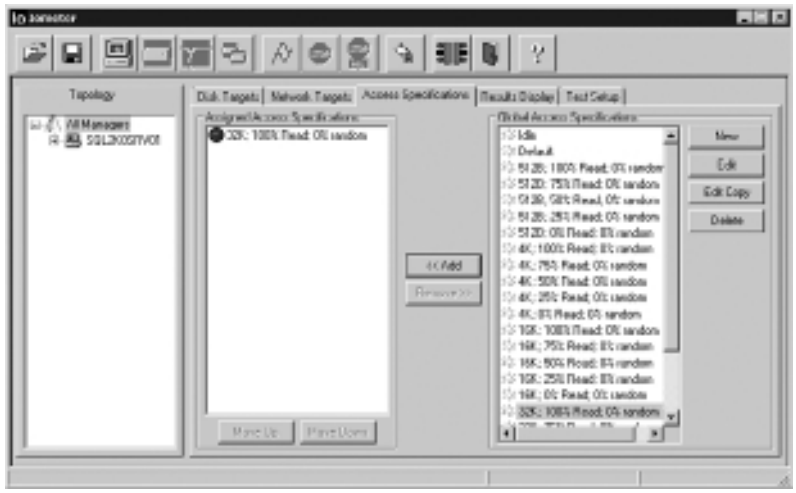

In addition to these generic I/O performance tools, you can also use the standard database benchmarks from the Transaction Processing Council (TPC) such as TPC-C, TPC-E, and TPC-H. These benchmark kits, however, require special skills to set up and run. Like any benchmark disclaimers, these performance tools will only help extrapolate your application's performance.

Remember that the tools we've covered will not represent actual SQL Server usage in your environment. You must also use realistic workloads that exercise SQL Server in the same way as used by front-end applications in your environment. Stress-testing the deployments on VMware vSphere with no "think times," for example, or comparing physical to virtual performance from different physical servers can be misleading. Evaluate the proposed workloads to gain an understanding of how current physical resource utilization peaks and heterogeneous workloads will affect resource utilization when these workloads are virtualized.

Test your current backup policies and procedures as appropriate. Ensure that the proposed consolidated environment will meet or exceed production requirements established by your SLAs if appropriate. Test your proposed availability scenario by forcing a failure and ensuring that you meet or exceed your target RPO/RTO objectives and that your data is safe and recoverable.

Monitoring Virtual SQL Server Performance

When monitoring your virtualized SQL Server application infrastructure, you have several options available to ensure that your SQL Server deployment on VMware vSphere platform is operating efficiently.

SQL Server DBAs typically use the Microsoft Windows performance-monitoring tool Perfmon to collect statistics. SQL Server integrates with Perfmon to provide familiar counters that indicate system performance. However, as with all measurement tools run in the guest operating system, time-based performance measurements are subject to error. The degree to which the measurements are inaccurate depends on the total load of the ESX host. Depending on the load placed

on the ESX host, the CPU and memory usage reported within the guest operating system can be different from what ESX reports.

You should use the VMware-provided tools:

- vCenter
- esxtop/resxtop

Figure 7.51 and Figure 7.52 provide examples for performance monitoring using vCenter and esxtop, respectively.

FIGURE 7.51
vCenter performance monitoring

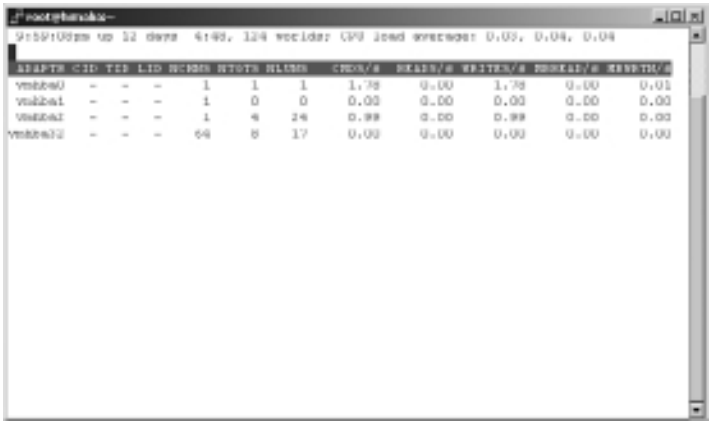

FIGURE 7.52
Performance monitoring using esxtop

Before we discuss specific performance counters to monitor, we'll address a question DBAs frequently ask how can they monitor SQL Server VMs in vSphere environment. In a consolidated environment with multiple virtual machines sharing hardware resources, DBAs are often concerned that their databases may not get the right amount of resources. vSphere maintains strict fairness in distributing resources across all running virtual machines based on resource

allocation policies. DBAs want to validate this on their own but cannot do it themselves since they typically do not have access to the previously mentioned vSphere monitoring tools. vSphere administrators usually control the access to these tools and for security reasons may not be willing to open up vSphere and vCenter access to other IT administrators. You can leverage the access control features in vCenter to work around this issue. To provide read-only access to select users, right-click the appropriate object in the vCenter inventory, and click Add Permission. In the Assign Permissions windows (Figure 7.53), you can add specific domain users and groups as appropriate with a read-only role.

FIGURE 7.53
vCenter read-only permissions to DBAs

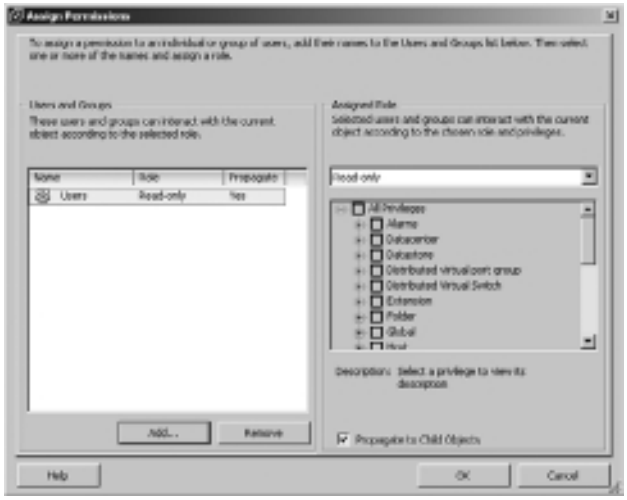

Users with the read-only role can view and report performance only for the objects they have access to; they cannot modify any objects in the vCenter inventory. All other options in the vCenter pop-up menus are disabled (see Figure 7.54).

FIGURE 7.54
vCenter read-only access

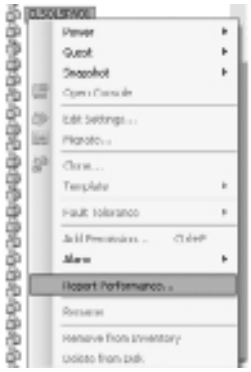

DBAs can leverage this read-only access to report performance for SQL Server VMs as well as the vSphere hosts and clusters where these VMs are running. You have access to the real-time

performance data from the Performance tab, or you can choose to export the performance data to an external file (see Figure 7.55).

FIGURE 7.55
Export performance data

When monitoring performance for virtual SQL Server deployments, pay close attention to the following VMware vSphere counters (Table 7.5), referring to online documentation on performance monitoring and analysis for more information on these counters and how to interpret them.

TABLE 7.5: VMware vSphere Performance Counters of Interest to SQL Administrators

SUBSYSTEM	ESXTOP COUNTER	VIRTUALCENTER COUNTER
CPU	%RDY	Ready
	%USED	Usage
Memory	%ACTV	Active
	SWW/s	Swapin
	SWR/s	Swapout
Storage	ACTV	Commands
	DAVG/cmd	deviceWriteLatency and deviceReadLatency
	KAVG/cmd	kernelWriteLatency and kernelReadLatency
Network	MbRX/s	packetsRx
	MbTX/s	packetsTx
	%DRPTX	droppedTx
	%DRPRX	droppedRx

Table 7.5 lists a few key counters you should add to your list of inspection points. Of the CPU counters, the total used time indicates system load. Ready time indicates overloaded CPU resources. A significant swap rate in the memory counters is a clear indication of a shortage of memory, and high device latencies in the storage section point to an overloaded or misconfigured array. Network traffic is not frequently the cause of SQL Server performance problems except when large amounts of iSCSI or NFS storage traffic are using a single network line. Check total throughput on the NICs to see whether the network is saturated on the ESX host.

Monitor I/O latency statistics and any disk I/O queuing. Use esxtop (or resxtop) to see reports of device latency, time spent in the kernel, and latency seen by the virtual machine. You can also monitor I/O queuing in the VMkernel by checking the number of queued commands. Figure 7.56 shows how you can monitor I/O latencies using the disk panel in esxtop. The figure shows the I/O latencies for devices connected to vmhba1 during a SQLIOSim test run in the proof-of-concept deployment.

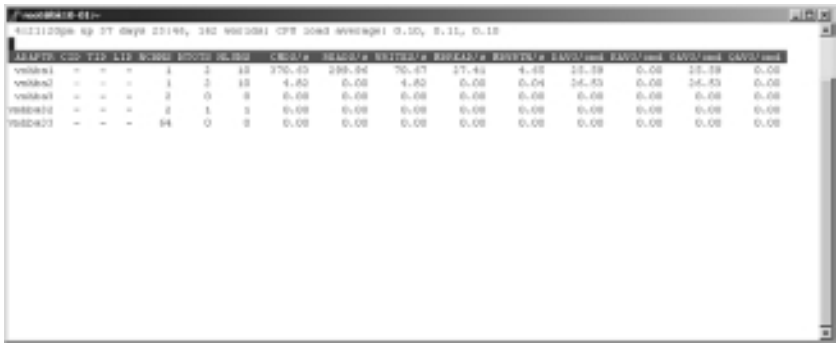

FIGURE 7.56 Monitoring I/O latencies using esxtop

`esxtop` also allows you break down the I/O latencies to specific storage targets. In Figure 7.57, you can see the I/O latencies experienced by the three virtual disks—to be used for data, transaction log, and TempDB—for the SQL Server virtual machine during a SQLIOSim test run.

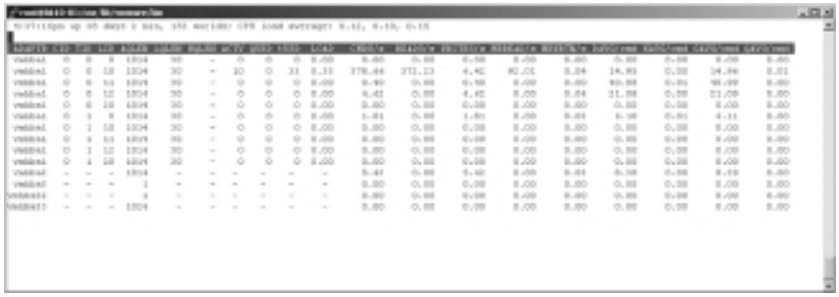

FIGURE 7.57 Performance monitoring using esxtop

You can leverage the customization feature in vCenter to monitor specific performance counters. Using custom chart settings, you can predefine the counters that vCenter displays. Figure 7.58 shows a custom chart setting to monitor storage-specific counters for SQL Server virtual machine in the proof-of-concept deployment.

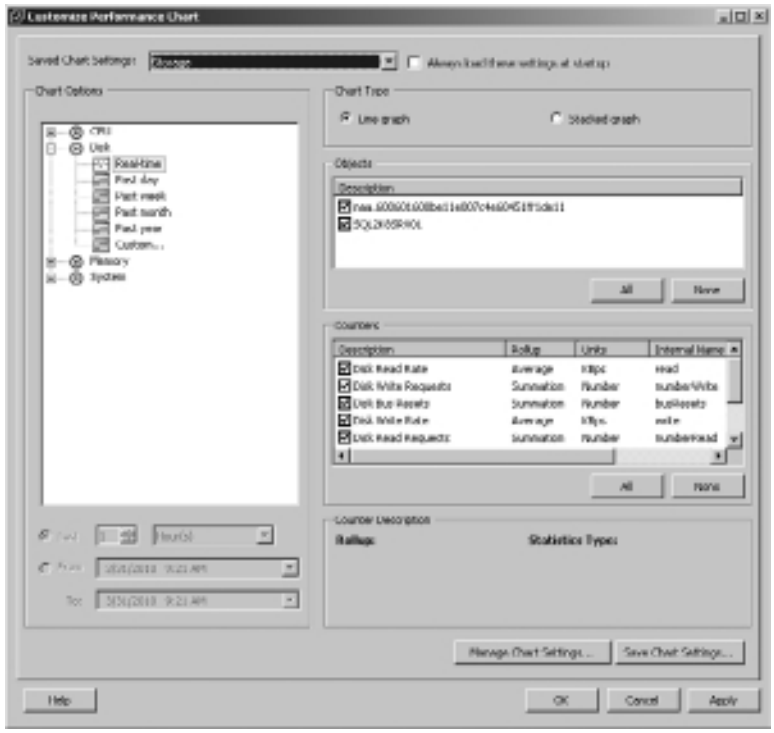

FIGURE 7.58
vCenter custom chart settings

In addition to esxtop and vCenter, vSphere also provides another tool, `vscsiStats`, to characterize virtual disk I/O workloads. `vscsiStats` collects performance data for storage activity at the virtual SCSI device level in the VMkernel so you can gather I/O information for specific virtual disks. You can start and stop online data collection for particular VMs and print various types of histograms for performance metrics such as `ioLength`, `seekDistance`, `outstandingIOs`, `latency`, and `interarrival`. We will provide you with a brief overview of how to use `vscsiStats` to gather this information for your virtual machine's virtual disks and then how the resulting data can be interpreted.

To use `vscsiStats`, you need service console access to the vSphere host where your virtual machine is running. `vscsiStats` is located in the folder `/usr/lib/vmware/bin`. You need to determine the worldGroupID (WIDs) of your virtual machine before you can collect data for the virtual machine. To do this, follow these steps:

1. Run `cd /usr/lib/vmware/bin` (press Enter).
2. Run `./vscsiStats -l` (case sensitive, with a lowercase *L*).

The output should list the WIDs of each of your virtual machines in an output similar to the following:

```
Virtual Machine worldGroupID: 5321, Virtual Machine Display Name: SQL2K8SRV01 {
 Virtual SCSI Disk handleID: 8216
```

```
 Virtual SCSI Disk handleID: 8217
 Virtual SCSI Disk handleID: 8218
 Virtual SCSI Disk handleID: 8219
}
```

This output shows that the SQL2K8SRV01 virtual machine has a worldGroupID of 5321. Next you use this ID to start monitoring virtual disks attached to this virtual machine.

```
./vscsiStats –s –w 5321
```

The output should look like the following:

```
[root@sd10-esx4-01 bin]# ./vscsiStats -s -w 5321
vscsiStats: Starting Vscsi stats collection for worldGroup 5321, handleID 8216
Success.
vscsiStats: Starting Vscsi stats collection for worldGroup 5321, handleID 8217
Success.
vscsiStats: Starting Vscsi stats collection for worldGroup 5321, handleID 8218
Success.
vscsiStats: Starting Vscsi stats collection for worldGroup 5321, handleID 8219
Success.
```

Notice that the specific virtual disks are identified by `handleID`s. You need these to track performance statistics for each individual virtual disk. Once `vscsiStats` starts collecting performance data, it can report various statistics such as `ioLength`, `seekDistance`, `outstandingIOs`, `latency`, and `interarrival`. You can view latency metrics for the previous SQL Server virtual machine by entering the following:

```
./vscsiStats -w 5321 -p latency
```

The output should look like the following:

```
Histogram: latency of Read IOs in Microseconds (us) for virtual machine
worldGroupID : 5321, virtual disk handleID : 8216 {
 min : 404
 max : 1728865
 mean : 38524
 count : 26146
 {
 0 (<= 1)
 0 (<= 10)
 0 (<= 100)
 17 (<= 500)
 508 (<= 1000)
 2534 (<= 5000)
 7524 (<= 15000)
 8547 (<= 30000)
 3751 (<= 50000)
 1596 (<= 100000)
 1669 (> 100000)
 }
}
```

The output provides minimum, maximum, average, count, and histogram distribution of latencies for all read I/Os for the virtual disk with a `handleID` setting of 8216. You can see that out of 26,146 read operations, about 90 percent completed with a latency of 50 milliseconds. You can use this information to optimize your storage configuration such as adding more drives to a LUN or trying faster disk drives.

This information can also be captured to a `.csv` file for use with a tool like Microsoft Excel. The command to output latency statistics to a file called `latency.csv` is as follows:

```
./vscsiStats -w 5321 -p latency –c > /tmp/latency.csv
```

To view the `seekDistance` metrics (the number of blocks between successive reads) for this virtual machine, you enter the following:

```
./vscsiStats -w 5321 -p seekDistance
```

`vscsiStats` should be used for the advanced analysis of workload I/O after you have configured and optimized storage configuration following the standard best practices and have performed regular monitoring using vCenter or esxtop.

VMware also provides an application management solution, vCenter AppSpeed, that can monitor the performance of SQL Server transactions. AppSpeed can detect the slowest or most used SQL transactions, as well as point out transaction dependencies across components in a multitier application stack. You can also set latency thresholds to enable ongoing performance service-level monitoring using AppSpeed (Figure 7.59). We have discussed the AppSpeed product in more detail in Chapter 5.

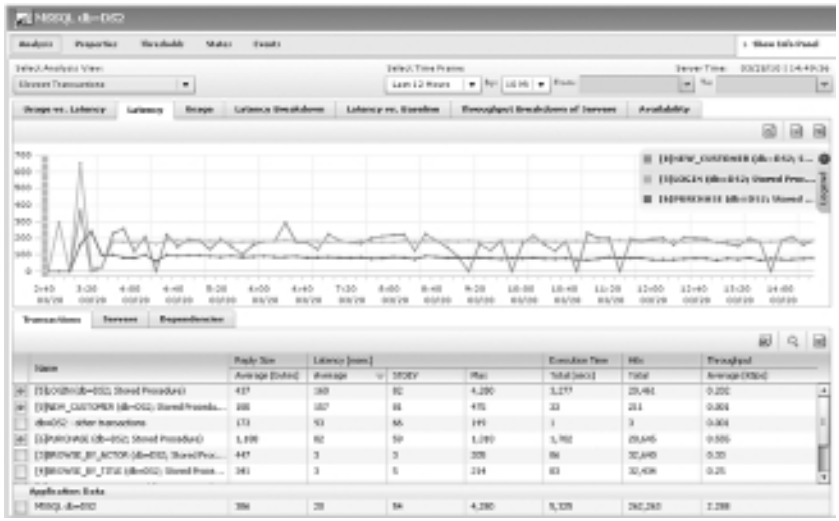

FIGURE 7.59 SQL Server monitoring using VMware AppSpeed

As stated earlier, you can use the scale-up approach where a single large virtual machine hosts multiple database schema inside a single SQL Server database or multiple databases in a single SQL Server instance. In either case, you can continue to use SQL Server Resource Governor to manage resource allocation across competing database workloads inside the SQL Server virtual machine. SQL Server Resource Governor provides the ability to define resource limits and priorities for different database workloads within a single SQL Server instance so

that runaway queries or large databases do not dominate the memory and processor resources. At the virtual machine level, resource allocation is controlled by the vSphere hypervisor and depends upon the virtual hardware configuration and DRS resource management settings.

When consolidating SQL Server in physical environments using multiple SQL Server instances, it is common to use Windows System Resource Manager (WSRM). When virtualizing SQL Server, we recommend that you create a separate SQL Server virtual machine per SQL Server instance, rather than running multiple SQL Server instances inside a single virtual machine. In this case, you would rather leverage VMware DRS to provide more efficient distribution of hardware resources among SQL Server virtual machines. DRS provides both proportional and absolute mechanisms to control priorities for various SQL Server workloads by using shares, reservations, and limits for CPU and memory; shares for virtual disks; and traffic shaping for the network. Also the resource allocation changes done using DRS take effect immediately without the need to restart the virtual machine. For example, in SQL Server environments virtualized using VMware vSphere, you can use these features to favor OLTP or DSS workloads as needed. During the batch or backup windows, you can set high shares for a specific database and low shares for OLTP workloads likely to be least used at that time.

Summary

In this chapter, we covered the considerations and best practices for virtualizing SQL Server using VMware vSphere. We discussed that SQL Server can be successfully virtualized without any data integrity issues. The vast majority of SQL Server deployments are good candidates for virtualization and can benefit from improvements offered by a virtualized infrastructure—such as improved management, availability, and scalability—thus reducing overall TCO.

Virtualizing SQL Server should provide benefits that go beyond server consolidation and lower total cost of ownership. Virtualization offers a more secure and flexible option to consolidate SQL Server compared to options available in a physical environment. A successful SQL Server deployment using VMware vSphere should provide better management and administration flexibility as well as higher levels of availability at a lower price. You can take advantage of such key VMware vSphere features as DRS, HA, and the SRM plug-in to eliminate planned downtime; reduce unplanned downtime; and recover rapidly from component, systematic, or site-wide outages. At the same time, you can continue to leverage the SQL Server high-availability features, such as failover clustering, database mirroring, and log shipping in virtual deployment.

To deploy SQL Server successfully in VMware vSphere, you must plan your virtualization project. First, you need to characterize your SQL Server workloads and clearly understand your availability, backup, and other operational needs. You must establish a baseline, using data from the currently running physical deployment. Using this baseline and other requirements, you can rank SQL Server instances to virtualize. It is critical that you follow best practices specific to Microsoft SQL Server and other vendor products in your environment as well as VMware vSphere best practices. In general, best practices in physical environments also apply to deployments on VMware vSphere without any changes. You must also determine your SQL Server placement strategy as well as methods to migrate SQL databases to the virtual environment.

Finally, we covered a proof-of-concept deployment of SQL Server using VMware vSphere that deployed a two-node failover cluster. Although your implementation will be different from our proof of concept, it is important to consider all of the best practices that we discussed in this chapter and showed how to implement in the proof of concept in order to ensure your SQL Server virtualization implementation is a successful one.

Chapter 8

Virtualizing Microsoft Office SharePoint Server 2007

In 2009 Microsoft estimated that SharePoint 2007 revenues were $5.6 billion, moving toward the $6.1 billion mark by fiscal year 2011. This signals widespread adoption of SharePoint over other similar technologies in the space. Microsoft SharePoint brings together several technologies and applications such as team collaboration, web content management, business intelligence, content management, records management, and publishing capabilities all in a single integrated application that provides an easy-to-use, consistent user interface. As the need for more green initiatives (such as reducing data-center space, power consumption, and cooling costs) increases and enterprises look to decrease capital expenses, more and more enterprises are virtualizing the Microsoft SharePoint application.

In this chapter, we will cover a bit of background information about SharePoint governance, SharePoint Server role responsibilities, key considerations for virtualizing SharePoint server roles, and the Microsoft and Hewlett-Packard (HP) SharePoint sizing tools. We will continue the data-center proof of concept by showing how to deploy a SharePoint farm for Source Discovery.

Why Do You Need SharePoint Governance?

SharePoint governance brings together a governing body that establishes the business objectives for your SharePoint design and deployment. SharePoint governance also defines the process and procedures regarding the operation, maintenance, and availability of the Microsoft SharePoint application. Governance brings together the people, processes, and procedures to ensure that the SharePoint design and deployment yields the most effective use, prevents information misuse and access, and ensures that the organization will achieve the best return on investment (ROI) long after the solution has been deployed into the organization.

Microsoft SharePoint technologies offer extreme flexibility and power in the areas of document control, records management, content management, and business intelligence. With this power and flexibility comes complexity. Given all the features and capabilities of SharePoint Server, you must carefully plan and design the solution. In addition, you must develop and then adhere to solid business objectives for SharePoint deployment. It is not uncommon for a well-designed and deployed SharePoint solution to be underutilized because of the complexity of the design or solution. Also, it's not uncommon for the organization to outgrow the intended design and implementation because of a lack of SharePoint governance. A well-defined and executed

SharePoint governance plan will ensure that the designed SharePoint solution is utilized for what it is intended and will prevent information access and misuse. So, what is SharePoint governance?

What Is SharePoint Governance?

Wikipedia defines governance as consistent management, cohesive policies, processes, procedures, and decision rights for a given area of responsibility. Governance is managing the deployment of information technologies. When you think of managing, you can use the project manager analogy. Project managers manage the costs, the risks, and the effectiveness of the solution once it is deployed.

Microsoft defines governance as a set of roles, responsibilities, and processes that you put in place within an enterprise to guide the development and use of a solution based upon SharePoint technologies. So, when we define *SharePoint governance*, we are defining governance with the objectives of managing the risks, adoption, and cost of the SharePoint solution in mind.

There are several roles and responsibilities within a SharePoint governance plan that will help ensure that you will deploy a SharePoint solution that meets the needs of the organization and that also possesses the qualities that will allow that implementation to grow and scale as the organization changes or evolves (Figure 8.1). We will touch on a few of these roles but will not attempt to explain in great detail.

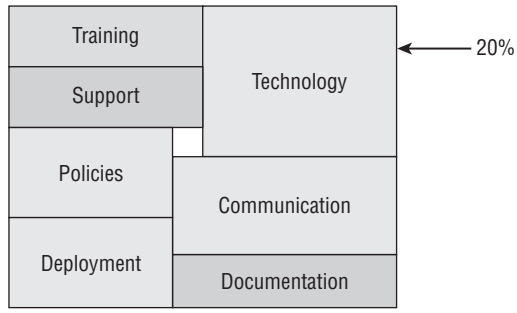

FIGURE 8.1 A simple SharePoint governance diagram

The following are just a few of the items that are part of a comprehensive SharePoint governance plan. In the following sections, we will cover one governance item from these six areas of a theoretical SharePoint governance plan.

- Project and operational management
- Development and configuration
- Infrastructure
- Operational processes and concerns
- Education and training
- Search, navigation, and taxonomy

TIP Microsoft provides a great deal of information about SharePoint governance and the roles and responsibilities associated with SharePoint governance. Consult the Microsoft Governance Resource Center for SharePoint Server 2007 website at `http://technet.microsoft.com/en-us/office/sharepointserver`, and click the Governance link for more information about SharePoint Server 2007 governance.

SharePoint Teams and Roles

The teams and roles area of SharePoint governance usually comes into play when undertaking the formal process of designing, deploying, and operating SharePoint in an organization (it therefore falls into the project and operational management category). Forming teams with specific roles and responsibilities can ensure that all aspects of a SharePoint implementation are properly governed. The teams are comprised of individuals who have specific roles and responsibilities for the solution (or solutions) that is based on SharePoint technologies. For example, the database administrator for a SharePoint implementation might be concerned with database design, scalability, high availability, and maintenance in a SharePoint implementation, while the web team and developers would be concerned with the SharePoint web parts and scalability of the web front ends. (Web parts define what a rendered SharePoint web page might look like.)

Customization Tools Used in the SharePoint Environment

This area falls under the category of development and configuration governance for SharePoint. Developers from different business units will want to have their own look and feel for their sites within the SharePoint environment. The challenge is to allow implementation while ensuring maintainability. For example, you want to ensure that consistent controls are used across the entire organization in order to permit the solution to be easily maintained while still allowing developers to create their unique look and feel for their sites. It is easy to create a solution that is extremely difficult to navigate and provides inefficient methods for accessing information.

SharePoint Environment Usage

This area of SharePoint governance falls under the infrastructure category. This area is more concerned with what SharePoint environments will be utilized for production, development, and test where code is to be implemented in SharePoint environments. The objectives for this area will be to ensure that code that is going to be used in the SharePoint environment is properly implemented through the SharePoint development and test environments before implementing it in the SharePoint production environment. In too many instances, code is rushed into production environments (other than just SharePoint) and causes irrevocable damage, which in turn causes lengthy recovery procedures and downtimes.

SharePoint Uptime and Downtime

This area of SharePoint governance is rooted in SharePoint operational processes and procedures. Here you focus mainly on defining service-level agreements, communication mechanisms for communicating service availability to end users, and processes and procedures for scheduled and unscheduled downtimes for the SharePoint application. Scheduled downtimes can encompass outages for site, database, and portal maintenance. Unplanned downtime could be the result of

service failures, database server outage, or any other event that causes the SharePoint application to become unavailable to end users. There usually will be well-defined processes and procedures to handle scheduled and unscheduled events in this area of governance.

SharePoint Community Training

In this area of SharePoint governance, you concern yourself with the topic of training the end users in the use of SharePoint applications. Although you can design a solution that can be 90 percent foolproof, oftentimes you can cover the other 10 percent of corner cases in which end users will have questions about simple to more advanced uses of the SharePoint application. This is the perfect opportunity often times to bring together the more advanced SharePoint users in community forums. In this space, the experts can effectively mentor the less knowledgeable users. Also, you can use brown-bag, one-on-one, and after-hours sessions to train and mentor inexperienced users, effectively reducing the demand on the more expensive human SharePoint resources and thus allowing them to operate more in strategic planning mode.

SharePoint Content Types

This area of SharePoint governance is derived from the navigation and search areas of governance and is more concerned with the core types of content that your SharePoint portal will process. Some of the more typical content types in SharePoint are Word documents (`.doc`), Excel spreadsheets (`.xls`), PowerPoint presentations (`.ppt`), and PDF files (`.pdf`). These content types can have searchable data or metadata associated with them to make searching and locating specific information much easier. When governing the content types, you will want to ensure consistency, which will make locating information easier. This can encouraged through the use of standardized fields.

SharePoint Server Roles and Responsibilities

SharePoint server technologies are created from a highly modular methodology, affording you the flexibility to design and implement in a variety of topologies starting from a single server deployment in which all server roles are part of a single server to a distributed model (more commonly known as *farm mode*). SharePoint server roles are really defined by the services that are running the server in the SharePoint farm.

You have a lot of flexibility when deploying in farm mode. You can combine server roles to reduce the number of servers required for your organization. For example, you could combine the web front-end server with a query so that you can reduce the number of hops that the information must travel between servers. Dedicating a web front-end server to do index crawling can also increase the efficiency of the crawling process as well as reduce the number of hops, thus improving the flow of data to the query servers. The following is a description of the roles:

Web front-end server The web front-end server role in the SharePoint server farm is responsible for responding to user requests for web pages. The web front-end server requests content from the SharePoint database and can also render the page (or pages) based on look and feel, cache the page, and return the relevant results to the end user. The web front-end server is highly scalable in that multiple web front ends can be deployed in a SharePoint farm

to provide aggregate bandwidth, load balancing, and redundancy in the event of single or multiple web front-end server failures.

Query server The query server role is responsible for responding to search requests from the web front-end role within the SharePoint server farm. Its secondary responsibility is to maintain a propagated copy of the index corpus that gets stored on the query server's local file system. Because the query role is one of the lightest SharePoint server roles, it is commonly combined with the SharePoint web front-end server role within the farm.

Application server This role in the SharePoint server farm consists of the additional applications that you may utilize in your server farm such as Excel Services, Calc Services, and document conversion services such as PDF and other documents types that are not available to out-of-the-box SharePoint Server implementations. These applications are typically combined onto a single server role.

Index server The index server role in a SharePoint server farm is responsible for maintaining an up-to-date index by crawling the index corpus. The *index corpus* is a collection of documents stored and crawled by SharePoint, including the categories and folders, document profiles, document versions that may exist, and role principals. A *role principal* is a unique Microsoft security identifier. A good example of role principals are Windows user accounts and group accounts.

NOTE If you have three Windows user accounts and four Windows group accounts, then you have seven role principals.

A single SharePoint Portal server can store about 1 million document versions and can crawl about 3.5 million documents. Because the index server can consume all of the memory resources of a server, you must carefully consider the virtualization of this role, as covered in the next sections.

Database back-end server The database role in the SharePoint server farm is responsible for storing, maintaining, and returning data to other roles in the SharePoint server farm. In typical SharePoint server farms, the database back-end server will host about 7 to 10 databases at a time. As your information architecture grows, you will need to plan for scaling out your database servers within the SharePoint server farm. Other considerations are high availability for your database back-end content because this is the "backbone" of your SharePoint server farm. Microsoft supports scaling out the SharePoint content database to multiple SQL database servers within a single SharePoint farm. Scaling out the content databases is relatively straightforward. Scaling out the other databases such as the admin, config, and share service provider (SSP) databases can prove challenging.

TIP Please refer to the Microsoft website for more information about scaling the SharePoint topology: `http://technet.microsoft.com/en-us/library/cc298578.aspx`.

Considering Server Role Virtualization Workload

Given the multiple server roles for SharePoint, you must consider each server workload characteristic individually. Some roles will be extremely disk I/O bound, whereas other roles will

be CPU bound. Some roles will place heavy demands on all of your server resources such as CPU, memory, network, and disk. Because of the diverse resource requirements for each server role, we will talk about each server role separately and the considerations for virtualizing each SharePoint server role.

You should properly plan your server role resources because any lack of the required resources for a server role will be felt throughout the SharePoint server farm via an unacceptable end user experience. We will also provide a brief view into the Microsoft and HP SharePoint sizing tools to help you get a feel for what your SharePoint farm will require in terms of server resources.

Web Front-End Server Role (and Query Role)

The SharePoint server's web front-end (WFE) role is a great SharePoint role for virtualization. This role is scalable by adding more web front ends to serve the user population. It is recommended that you combine the query server role with this role to reduce the number of virtual machines you would normally deploy to support the query role.

The query role is a very lightweight role in that, on average, it barely requires any resources to perform effectively. Because the WFE role is highly scalable and SharePoint currently uses a server licensing model, it is recommended that you carefully consider the number of WFE server roles in your SharePoint implementation. The VMware Solutions group (as well as other partners) has performed scalability testing of the WFE roles in a virtualized SharePoint solution and found that a virtual machine with a single virtual CPU (vCPU) and 2 GB of memory will slightly outperform a two-vCPU virtual machine with the same 2 GB of memory.

TIP You can find the results of these tests on the VMware website under Solutions: `www.vmware.com/solutions/business-critical-apps/sharepoint_server/resources.html`.

When deploying the WFE role in your SharePoint environment, you must consider the ratio of WFE servers to the number of SQL database servers. Also, do not deploy all of your WFE virtual machines onto the same ESX host in order to prevent single points of failure. The use of DRS anti-affinity rules can prevent the WFE virtual machines from migrating to the same ESX host. If you are deploying more than two WFE virtual machines in your SharePoint implementation, consider using a hardware load-balancing solution because it can reduce the demand on the WFE virtual machines.

Consider cost and performance when deciding whether to use hardware load balancing vs. software load balancing. If you are deploying software load balancing, deploy dedicated virtual network adapters that are connected to two virtual switches (vSwitches), which in turn should be connected to at least two physical network adapters for the public and load-balancing traffic. Microsoft recommends that you deploy no more than eight WFE servers to every SQL database server and no more than three web servers per single domain controller in your SharePoint server farm to prevent degradation to SQL Server performance. If you have Microsoft "all you can eat licensing," we recommend deploying your WFE virtual machine roles with a single virtual CPU and 2 GB of memory. For good overall performance, the two-vCPU virtual machine with 2 GB of memory also performs great and will suit more than 98 percent of the needs of those deploying the SharePoint WFE server role.

To facilitate ease of deployment of this role, we recommend that you use a template with all of the SharePoint prerequisites such as Microsoft Internet Information Services (IIS) and

.NET Framework versions v2.0 and v3.0 installed on the template. Microsoft recommends that the query catalog size be at least 30 percent of the SharePoint indexable content. The VMware Solutions team generally sees about 10 to 15 percent. You can also improve performance by disabling IIS logging if it is not required. IIS logging captures all the information about web functions including errors, warnings, and informational events on the server.

You will also need to configure three virtual disks:

- One for the operating system
- One for the installation of the application binaries
- One for the query LUN to support the SharePoint query role

These disks can be VMware File System (VMFS) disks. We will cover the creation and configuration of the WFE role virtual machine in more detail in the "SharePoint Proof of Concept" section of this chapter. You can also locate the SharePoint binaries on the template to eliminate the need to mount a compact disc or an ISO image to install the SharePoint application.

NOTE Tests have shown that you may not realize significant throughput gains after deploying more than five web servers.

Application Server Role

The SharePoint application server role is similar in workload requirements to the WFE role. Most of the same rules will apply to this role, with the exception of multiple network adapters because this role cannot be load balanced. The disk configuration can also reduced to two virtual machine disks (VMDKs)—one for the operating system disk and the other for holding the installed binaries for the SharePoint installation programs. As your resource requirements for your application services increase, you can split up application services onto addition application server virtual machines in the SharePoint farm.

Index Server Role

The index server role can be a tricky role to virtualize because the performance will greatly depend on the size of the stored content or corpus to be indexed. Indexing is very demanding on CPU, memory, and disk I/O resources. Carefully consider the size of content to be indexed. For small to medium SharePoint farms of less than 1,000 users, we recommend that you start with a two-vCPU virtual machine with 4 GB of memory and increase virtual CPU and memory as needed. For large SharePoint farms, we recommend that you start with a four-vCPU virtual machine with 6 GB memory and increase memory and vCPU as needed. Utilize separate virtual disks for your operating system and index corpus, and use the fastest disk possible.

When creating your operating system virtual disk, ensure that you allocate enough space to hold the SharePoint and IIS logs. To reduce the number of network hops, use a dedicated WFE for crawling the index content. This can be done by enabling the WFE service on the index role. For best performance, use a thick-provisioned virtual disk that is carved out of separate disk spindles on the storage array because this will prevent any contention among other disks that are carved out of the storage array.

Database Back-End Server Role

The SharePoint database back-end server role is another tricky role to virtualize. This role will exhibit the highest CPU, memory, and disk I/O requirements of all the SharePoint server roles. Because some of the SharePoint operations don't present SQL-friendly statements to the SQL Server database engine, some SharePoint operations will place heavy demands on all of the host server's resources. These are some of the operations that can cause high demand on the SQL Server:

- Document uploads
- Document check-in/checkout
- Full index crawls
- Movement of documents

It is important that you follow the Microsoft guidelines and recommendations for deploying the SharePoint database back-end server role. Some considerations when deploying this role are to follow the Microsoft best practices for distributing content databases and keeping the site collection no larger than 100 GB to prevent SQL Server performance degradation. Use SQL Server aliases for your SharePoint SQL Server instances. For larger SharePoint deployments, consider deploying a four-vCPU virtual machine with 8 GB of memory, and increase vCPU and memory resources as needed. If you are running a four-vCPU virtual machine and find that increasing the vCPU resources to an eight-vCPU virtual machine is still unacceptable (which we doubt), you should consider scaling out the SharePoint database server back-end role onto multiple SQL Server instances.

Microsoft supports SharePoint/SQL Server scale-out scenarios. Because the SQL Server will host about 7 to 11 databases, you will want to ensure that you plan for adequate disk space for your SQL Server database back end. For larger SharePoint environments with heavy content modification, check-in/checkout, search queries that return a large amount of hits, and deep document versioning, we recommend you use a RAID 10 disk configuration for the SharePoint and TempDB databases and RAID 10 or RAID 1 for the SharePoint and TempDB transaction log files.

You should monitor your SharePoint implementation on a regular basis to ensure that your implementation is meeting or exceeding the service-level agreements established by your SharePoint governance board. Another periodic task that should be performed on your SharePoint farm is to execute a database defragmentation maintenance plan to ensure that your SharePoint database back-end server(s) are performing optimally. These things should be covered under your SharePoint governance plan and carried out by the appropriate personnel designated by the governance board for SharePoint.

NOTE You can find more information regarding the defragmentation of the SharePoint databases on the Microsoft website at `http://support.microsoft.com/kb/943345`.

Sizing the Server Roles

In the previous section, we talked about the SharePoint server role workload characteristics. In this section, we will use the HP and Microsoft SharePoint sizing tools to generate a sample configuration for a SharePoint farm that includes high availability. Although these sizing tools do a great job at sizing physical servers, in a virtualization solution that uses very powerful or "big-iron" hardware, there will be the opportunity to reduce a lot of server and storage requirements

generated by these tools. The guidance produced by these tools should serve only as a baseline data point. You should thoroughly understand your information architecture (content needs, locality of data, content types, indexing needs, and so on). With that said, we'll start by covering the Microsoft System Capacity Planner 2007 tool for SharePoint. In the examples, we will show how to create a sample SharePoint farm configuration. In our sizing scenario for both tools, we will create a five-server SharePoint farm with the following attributes:

- 100 users with average collaboration
- High availability for the web front-end servers
- High availability for the database back end
- 15,000 RPM SAN-based storage, 300 GB enterprise Fibre Channel disks

Using the Microsoft System Center Capacity Planner 2007

We will cover some of the basic steps for navigating the tools and creating a sample SharePoint configuration for a small to medium business (SMB) environment. Again, as we mentioned earlier, when you create these configurations for your SharePoint environment, you should use these sample configurations as a baseline data point in your design. Virtualization will help consolidate and reduce the need for some of hardware requirements that these tools are suggesting. The first step is to download Microsoft System Center Capacity Planner 2007 from the Microsoft website (`http://technet.microsoft.com/en-us/systemcenter/scx/bb969059.aspx`) and install it.

NOTE The Microsoft System Center Capacity Planner 2007 requires that Microsoft System Center 2007 v2.0 or newer be installed before installing the tool.

Getting the Basic Information

To generate a sample configuration for SharePoint from the System Center Capacity Planner for SharePoint, you must provide the tool with basic information about your proposed SharePoint configuration such as the following:

- Name of the SharePoint farm
- Logical topology (Internet or intranet) and usage profile (light/average/heavy)
- Number of users (local and any branch offices)
- High-availability requirements (front end and back ends)
- Storage resource requirements

Opening the System Center Capacity Planner Tool

To use the Capacity Planner tool, perform the following steps:

1. Select Start ➢ SharePoint Capacity Planning Tool, and select the SharePoint Capacity Model application.
2. The SharePoint System Center Capacity Planner Tool will open. Select Microsoft Office SharePoint Server 2007 in the drop-down menu under Capacity Models (Figure 8.2).

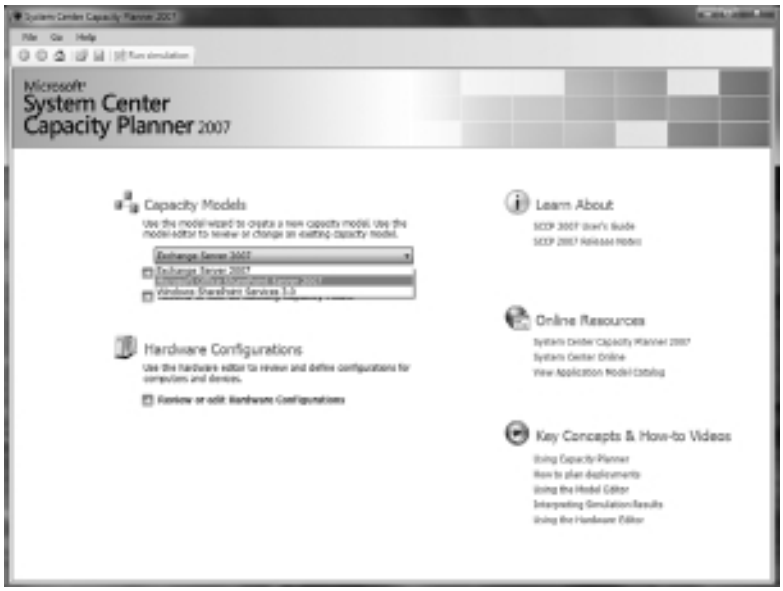

FIGURE 8.2
Selecting Microsoft Office SharePoint Server 2007 in the drop-down

After you have made the selection, the capacity model for SharePoint will load.

3. Click the Create A New Capacity Model link, which will load the Model Wizard for creating your capacity model.
4. Once in the Model Wizard, enter the following information:
   - The name of the SharePoint farm.
   - The number of users/clients the SharePoint farm will host.
   - The type of deployment (Internet or intranet).
   - The usage profile (there are many to select from). For this example, select the Average Collaboration profile, as shown in Figure 8.3.

There is also an option that allows you to enter custom values for your usage profile. This is where you can adjust the different SharePoint activities to more finely tune your workload requirements within the tools. You can do the following:

- Adjust actions such as home page access
- Download documents
- Search operations
- List pages
- List transactions per hour for connected clients (Figure 8.4)

FIGURE 8.3
Model Wizard screen for entering SharePoint basic information

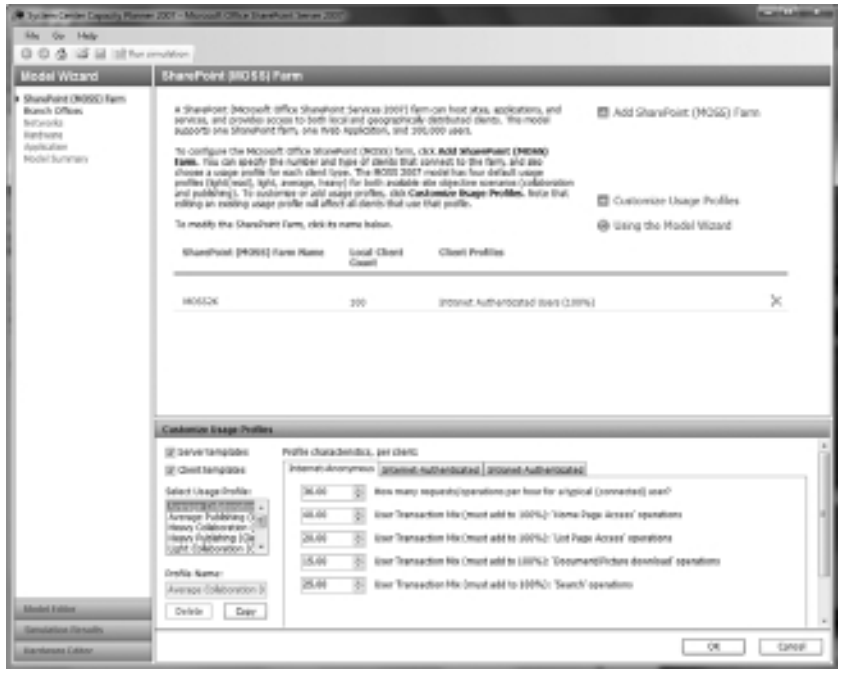

FIGURE 8.4
Customizing the SharePoint workload transaction mix

5. Click Next, and you will be taken to the Networks page (Figure 8.5). Unless you have branch offices, you will not have the option of changing anything on this page. If you do define branch offices, you will given the opportunity to edit the parameters for each branch office such as number of users, usage profile, and network bandwidth between sites.

FIGURE 8.5
Networks configuration page

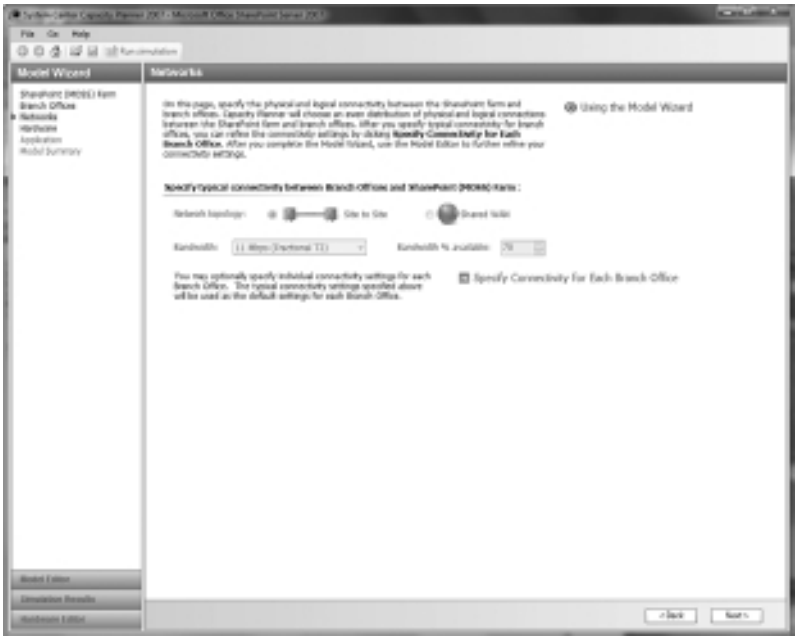

6. Click the Next button, and you will be taken to the Hardware page of the Model Wizard. This is where you will define your server hardware and storage requirements for your SharePoint farm. For server hardware, you will have the option of selecting from Intel and AMD multicore processors and various sizes, types, and speeds of disks ranging in size from 32 GB to 1 TB and ranging in speeds from 7,200 to 15,000 RPM as well as SATA, SAS, and SCSI disks. In the disk section of the Hardware page, you select the type of disk, speed, and size for each SharePoint server role (Figure 8.6).
7. Click the Next button, and you will be taken to the application requirements page where you will be asked about high availability for your SharePoint configuration (Figure 8.7). You can choose from various types of availability:
   - Failover clustering
   - Database mirroring
   - SQL Server log shipping

These are all SQL Server high-availability options for SharePoint. On this same page, you will also be asked about the size of the content required for your SharePoint content databases.

After you have defined your content database size and your high availability requirements, click the Next button and continue to the model summary.

TIP Before choosing a high-availability model, check with your SharePoint governance board to ensure that your chosen high-availability model will meet your requirements.

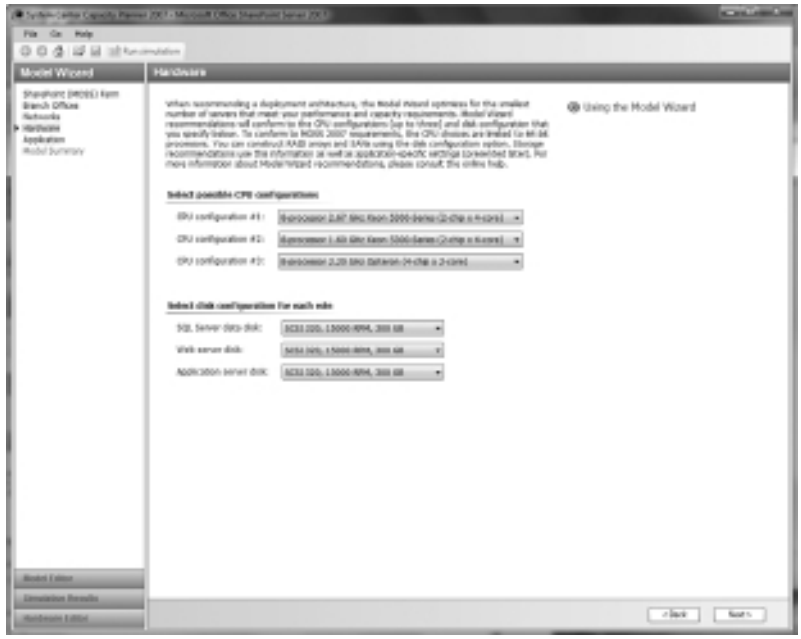

FIGURE 8.6
Selecting the server and disk hardware

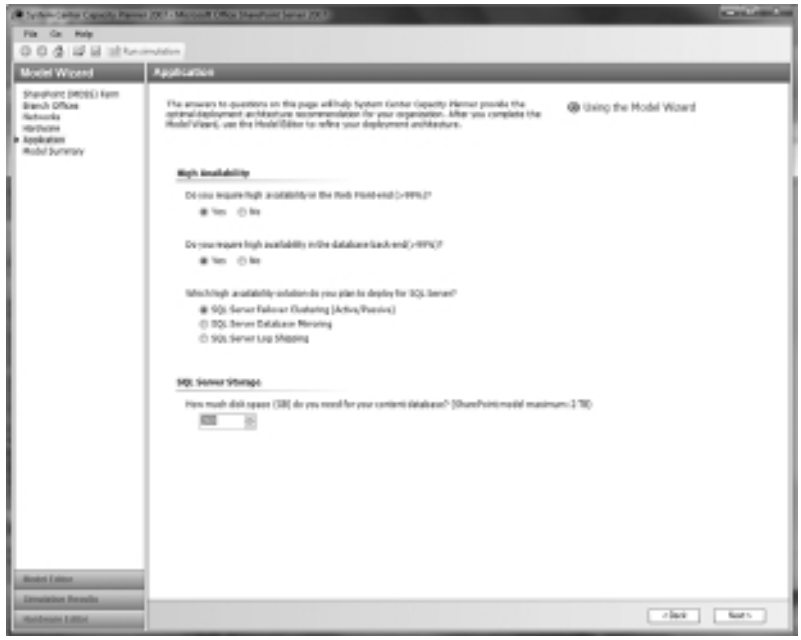

FIGURE 8.7
SharePoint Application configuration page

8. The Summary page lists all of your requirements for your SharePoint configuration mapped to the chosen hardware and SharePoint server roles.

NOTE At the top of the model summary, you will notice that there is a message about the topology not taking into account existing wide area network (WAN) utilization. In your planning, ensure that you know this information if you are going to have a distributed SharePoint implementation.

If you have chosen to provide high availability for your back-end databases, you will see a SQL cluster configuration for your SharePoint SQL Server configuration as well as multiple web front ends if you have chosen to implement availability for the front end. Here you will have the option of using hardware or software load balancing to provide redundancy and availability for your SharePoint web front-end server roles (Figure 8.8).

FIGURE 8.8 Summary for the SharePoint capacity model

At the top of the model's Summary page, click the Run Simulation button to test your configuration and see whether your configuration will meet the requirements set forth by your SharePoint governance body. Once the simulation has completed, you will see a summary of the results for your SharePoint capacity model (Figure 8.9). Again, as we mentioned earlier in this chapter, use this guidance only as a baseline data point with other information about your SharePoint requirements.

FIGURE 8.9
SharePoint capacity model simulation results page

Be sure to carefully review all information on the results page. Some warnings generated for actions might not be cause for concern, such as moving massive amounts of documents from one workspace to another. This is an example of an action that does not often occur, so it would generally not be a cause for concern.

If you click the Model Editor link, you will get a visual view of your selected SharePoint topology (Figure 8.10). On this page you can edit and modify the components of the SharePoint topology. You can add and remove servers, networks, and other items to improve the configuration and simulation performance. There are also links to help you on each page of the Model Wizard. Once you are satisfied with your configuration, you can save the model and close the tool. You can load the model for use again at a later time if needed.

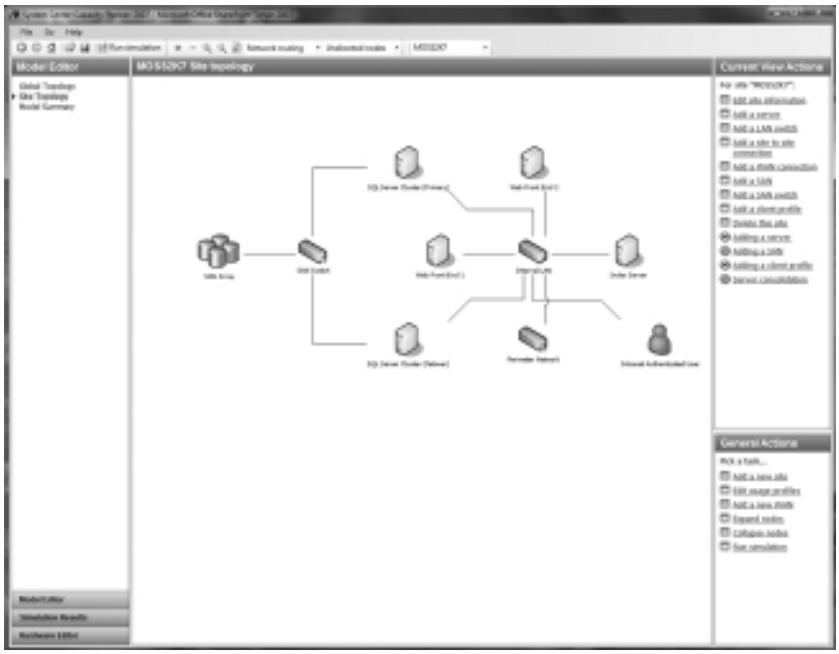

FIGURE 8.10
Visual view of the chosen SharePoint topology

Using the HP Sizing Tool

The HP sizing tool for Microsoft Office SharePoint 2007 was developed by HP for modeling SharePoint configurations and solutions based on HP ProLiant hardware and StorageWorks SAN arrays and direct-attached storage technologies from HP. The tool can generate SharePoint configurations as well as packaged solutions from HP. Here we'll show you how to generate a SharePoint configuration. We will show you the pricing information as well. The tool is freely downloadable from the HP website and can be installed by simply running the installer. The tool will periodically update itself from the HP website. Follow this procedure to use the tool:

1. When you first launch the tool, you will be asked to fill out the profile information for your company. Fill out the name, company, and email information to continue to the tool (Figure 8.11). Your name and email are required and must be entered before you can continue.

FIGURE 8.11
Entering your profile information

2. Click OK to continue, and you will be presented with the Sizer Home options screen.

3. On this screen, click the Build A Solution option, and you will be taken to introduction page (Figure 8.12). From here you will be able to customize your SharePoint solution in great detail.

FIGURE 8.12
Build A Solution introduction

4. Click the Next button to continue building and customizing your SharePoint configuration.
5. On the Solution Reference page, enter your customer or project information and any details about the project for future reference (Figure 8.13). You can save projects to refer to a generated configuration/solution as a guide or rule-of-thumb configuration for other projects in the future. Click the Next button to continue.
6. You are now on the Business Profile page (Figure 8.14). Choose the business profile for which you are designing the SharePoint configuration/solution. On this page, you get to choose from the three profiles:
   - SMB: 100 users or less
   - Mid-market: Up to 1000 users
   - Enterprise: 1,000 users or more

Using these three models, you should be able to design a solution that will fit into 99.99 percent of use cases.

FIGURE 8.13
Solution Reference page

FIGURE 8.14
Business Profile page

7. For the sample project, choose the SMB profile, because you will be designing the sample with 100 users or less. Click the Next button to continue with the design of the sample project.

8. Now you have arrived on the Usage Data page (Figure 8.15). On this page, define how the solution will be utilized:

   - The exact number of users who will access the SharePoint solution
   - The concurrency at which the users will access the solution

   When we talk about concurrency, we are defining the number of users who can access the solution simultaneously. We like to design our solution concurrency in the 90 percentile range because this gives a system a certain amount of headroom; in most cases, all of your employees who can access the solution will not log into and use the solution at exactly the same time.

   Click the Next button.

FIGURE 8.15 Selecting the usage pattern for the SharePoint solution

9. You are now at the Solution Component page (Figure 8.16). On this page you have the opportunity to customize the workload mix. In most SharePoint design scenarios, you would go with an out-of-the-box (OOTB) template and deploy based on that. Here you

actually get to customize your design by modifying the default values for the template. The following are some of the workload actions that can be modified:

- Portal usage percentage
- Search function percentage
- Team function percentage
- Level of content modification

Using these workload action parameters, you can tune your solution to model your usage within your organization or to the operational metrics your governance board has agreed upon. Click the Next button to continue designing the sample solution.

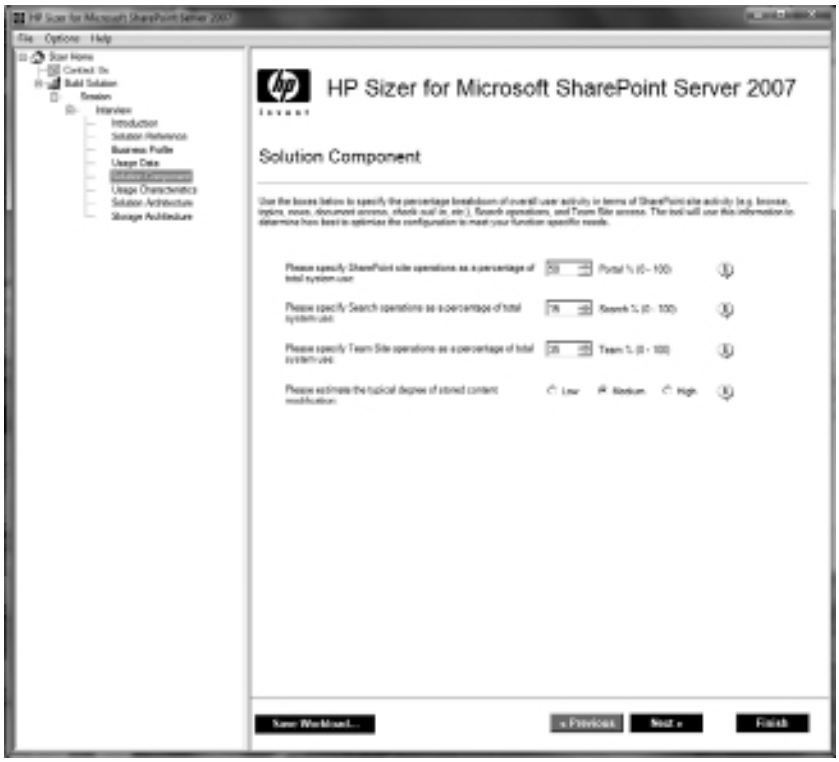

FIGURE 8.16 Customizing the SharePoint workload mix on the Solution Component page

10. Now you're on Usage Characteristics page where you specify how you will use the solution (Figure 8.17). You will be determining the following:
    - Whether you need high availability for your solution
    - Whether you will doing heavy indexing of the content data
    - The size of the index catalog

- How much disk space you want to allow the solution to use
- How large the content database is
- Other relevant usage characteristics that will have direct bearing on the performance of the solution

Click the Next button to continue to the Solution Architecture page.

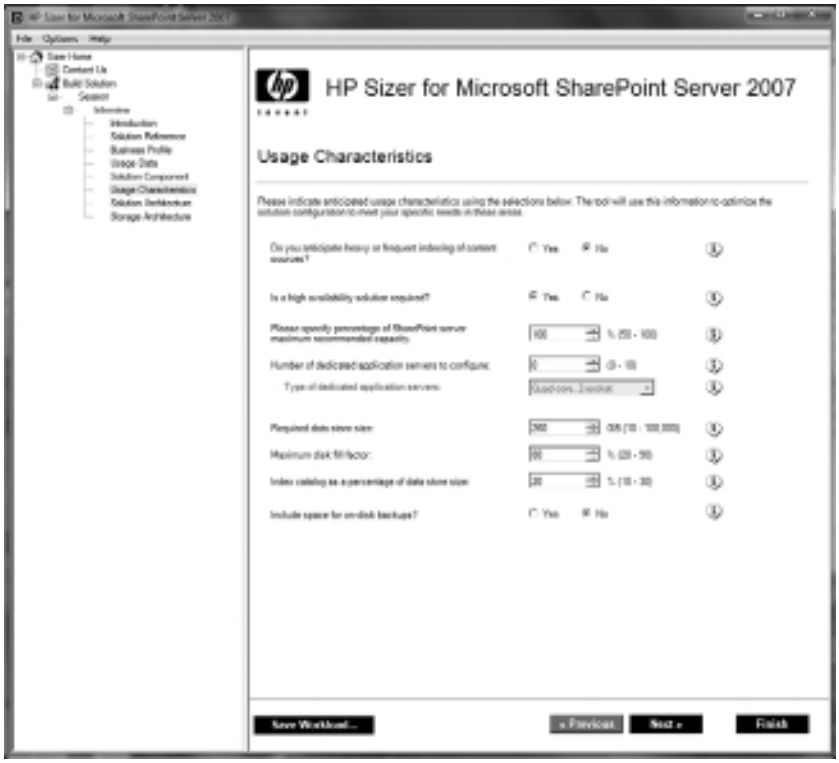

FIGURE 8.17 Defining the SharePoint solution's usage characteristics

11. On the Solution Architecture page, select the server architecture and footprint for your solution (Figure 8.18). Here you can choose the following:
    - AMD or Intel server architecture for your SharePoint solution.
    - The blade server or rack-mount server footprint.
    - The operating system and SQL server version for your solution. We recommend using Windows Server 2008 64-bit edition and SQL Server 2008 for your SharePoint configuration.

Click the Next button to continue.

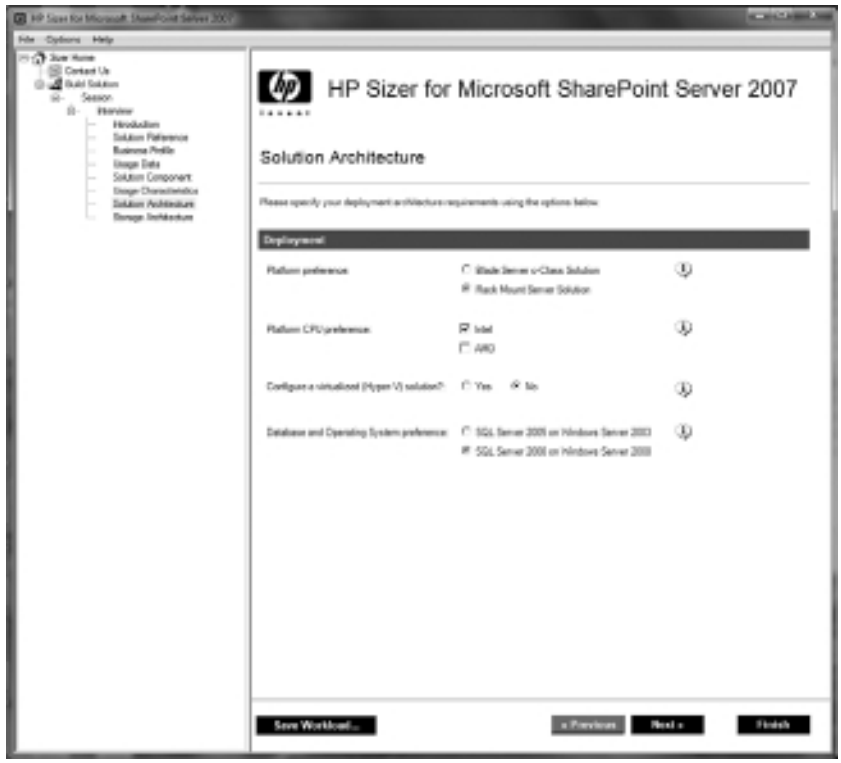

FIGURE 8.18 Choosing the SharePoint Server architecture

12. This takes you to the Storage Architecture page, so you are coming to the end of configuration of the SharePoint solution (Figure 8.19). On this page you will choose the storage solution for your SharePoint solution. You can allow the sizing tool to choose your storage for you, or you can choose to customize the storage you intend to deploy. Your choices are as follows:

    - Direct-attached storage (DAS)
    - SAN storage

    In some cases you might be able to select SATA-based disks depending on your workload. The sizing tool is smart enough to disable this option based on your projected workload characteristics.

    Click the Finish button to see the configuration details of your SharePoint solution.

FIGURE 8.19
Selecting the storage architecture for your SharePoint solution

The SharePoint solution configuration details are presented (Figure 8.20). You'll see configuration details such as the following:

- The number of servers
- The type of disks
- The number of disks needed for the solution
- The configuration for each server role
- The network bandwidth requirements for the solution

Understand that the solution generated is all based on HP technologies and pricing. Using virtualization will allow you to consolidate some of these roles onto larger, more capable hardware; however, it's good to get a view into what kind of resources would be required for deploying different SharePoint solutions.

FIGURE 8.20 SharePoint solution configuration result details

If you click the Calculate Full Solution button at the bottom of the Solution Lookahead pop-up window, you will get the HP pricing information for the configured solution that you have created (Figure 8.21). The tool will periodically update the pricing files and attempt to maintain the current HP pricing for the servers and storage in your selected configuration. Other options on this page allow you to save on the solution pricing, further customize the solution, and get more solution details.

In this section, we covered two of more popular SharePoint sizing tools and their options for configuring and sizing a SharePoint Server farm solution. The Microsoft sizing tool is totally agnostic and does not suggest any specific vendor's server hardware and storage for configuration. The HP sizing tool, on the other hand, designs its SharePoint sizing and configurations based on HP's server hardware and storage technologies.

As we mentioned earlier, virtualization on "big-iron" servers will allow you to consolidate and reduce the amount of physical server hardware that you have to purchase and deploy and will also reduce the server provisioning time. Use the sizing configuration information as a data point in your SharePoint design, and communicate the configurations and design to your governance board to ensure the solution fits the governance board's requirements for using SharePoint in your organization.

FIGURE 8.21
Viewing the Solution Alternatives page

SharePoint Proof of Concept

Source Discovery has experienced great success in its virtualization projects. It has virtualized critical Windows Server instances, Active Directory, IIS, and SQL Server. These projects have had great success and have transitioned into the production environment. Currently, Source Discovery is using a team collaboration application that is currently reaching its end-of-life serviceability. Source Discovery has investigated the possibility of using Microsoft Office SharePoint Server as a replacement for this application.

It has formed a governance body and included all of the necessary stakeholders in the design process. It has been decided that the company will do the following:

- Deploy a SharePoint server farm consisting of six servers. Because this proof-of-concept SharePoint server farm will host production data, the company will require backups of this farm.
- The company will use a third-party product it has licensed and already has in-house.
- The company will perform full and differential backups on the farm.
- The company will also require high availability using Microsoft Failover Clustering for the back-end database servers and a dedicated index server.

During the proof of concept, the company will use the Windows Network Load Balancing (WNLB) features of the Windows Server operating system to provide failover and redundancy for the two web front-end servers. When Source Discovery transitions the farm to production, it will use a hardware load balancer to provide load distribution, redundancy, and availability for the SharePoint farm. The company has decided to virtualize all aspects of SharePoint Server 2007 for the proof of concept. Using the SQL Server failover cluster examples provided in Chapters 3 and 7, we will show how to build a clustered SQL Server configuration for the SharePoint Server farm.

We will show how to use templates from Chapter 3 to build the SharePoint servers for the index and web front-end servers. As you deploy these servers, you will configure the vCPU, memory, disks, and networking appropriately to match the application's requirements. (We are assuming that you have a thorough knowledge of the Microsoft Office SharePoint Server 2007 product because we will not cover the configuration of the product in this proof of concept.) In the next sections, we will cover the configuration for the servers in the farm. We will start with the database back end and work our way to the front-end web servers.

Prior to creating any of the virtual machines, we created several the SharePoint Server 2007 administrative service accounts. You could have all of your services use one account but this would not be a best practice for deployment of this application .Evaluate the security requirements of your organization when considering account provisioning for SharePoint Server.

TIP For more information on SharePoint administrative service accounts, please click the Plan For Administrative And Service Accounts link at `http://technet.microsoft.com/en-us/library/cc263445.aspx`.

Building and Configuring the SQL Server Database Back-End Virtual Machine

In Chapter 3, you learned how to build a virtualized Windows failover cluster. Using the vSphere platform, you have these options when building Microsoft clusters:

Cluster in a box (CIB) This is a Microsoft cluster in which all nodes of the cluster reside on the same ESX host and use `.vmdk` disks as shared storage. All networking is handled through the single ESX host.

Cluster across box (CAB) A Microsoft cluster places the nodes of the cluster on different ESX hosts. Networking is handled through dedicated vSwitches on each ESX host. Storage is provided using raw device mapping (RDM) disks mounted on a virtual disk controller that is configured to share its disks across ESX hosts.

Physical to virtual In this configuration, a virtual machine cluster node is added to the physical cluster. The networking is handled through the vSwitch on the ESX host on which the clustered virtual machine is residing. Storage is handled in the same way as in the CAB configuration in which RDMs are used to provide the storage used in the physical cluster to the virtual cluster node.

For this example, the cluster configuration for the SharePoint Server farm database back end is a Microsoft failover cluster that is configured as a CAB. We're using SQL Server 2008 Enterprise edition to provide scalability. Following all of the guidance provided in Chapters 3

and 7, we have a functioning cluster. Figure 8.22 shows the configuration of the SQL Server virtual machine.

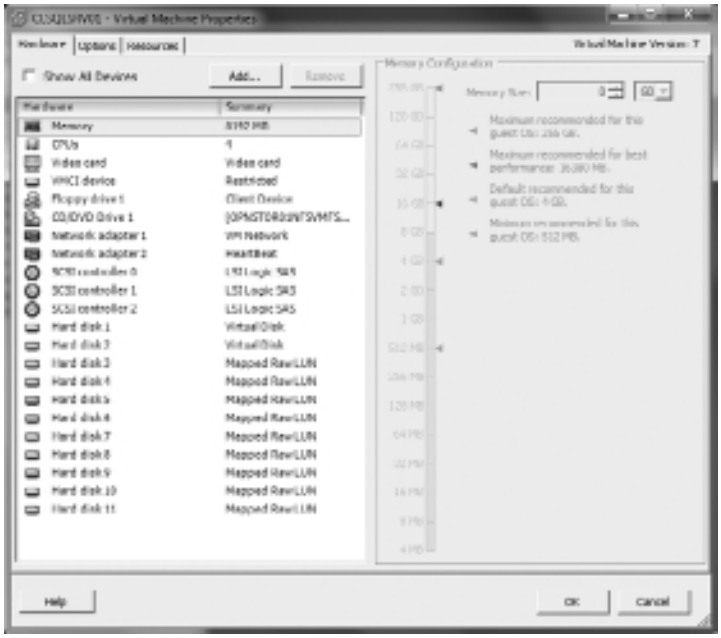

FIGURE 8.22
Clustered-node virtual machine

In this figure, you can see that we have provided the virtual machine with the following hardware:

- Four vCPUs.
- Eight GB of virtual machine memory.
- Two virtual network interface cards (vNICs): a public network interface for client access and cluster management and a heartbeat network interface for cluster traffic. This heartbeat network is used to provide cluster health information to each node in the cluster and to ensure that there is timely communication between cluster nodes to prevent false-positive cluster failovers and split-brain cluster scenarios.
- Three vSAS disk controllers. Separate disk controllers are required for the operating system disks and the clustered disks. Here we have employed a third vSAS disk controller to separate the database disks from the transaction log disks. This will provide more I/O throughput from the virtual machine as well as provide the ability to separate the database I/O from the transaction log I/O. Databases use random I/O patterns, whereas transaction log I/O uses sequential I/O patterns. Separating these I/O patterns results in more efficient utilization of the storage array.

This configuration is duplicated on different ESX hosts.

The cluster is then configured using the Windows Failover Cluster Management utility provided when the Windows failover cluster role is installed onto the Windows operating system.

Figure 8.23 shows the final cluster we have built in preparation for the clustered SQL Server database back end.

FIGURE 8.23
The Windows failover cluster

In this figure, we created the SQLCluster.sourdisc.com Windows failover cluster. As you can see, there are two nodes in this failover cluster. You will notice that there is also a SQLClusterDtc resource in the figure. This is the distributed transaction coordinator (DTC) resource required for a clustered SQL Server configuration. The SQL Server software will not install in a clustered configuration without the creation of the clustered DTC resource. Using the clustered events log, you can see any warnings and errors that might prevent the cluster from performing optimally.

Once you configure the Windows failover cluster with the clustered DTC resource, you install the SQL Server 2008 Enterprise edition in a clustered configuration. This is a two-part installation:

1. Install the software on the first node of the Windows failover cluster to create a single-node SQL failover cluster.
2. Repeat the install, and add the second node of the SQL Server failover cluster.

Figure 8.24 shows how to select the SQL Server cluster configuration.

SQL Server Installation Center
Planning
Installation
Maintenance
Tools
Resources
Advanced
Options
New SQL Server stand-alone installation or add features to an existing installation
Launch a wizard to install SQL Server 2008 in a non-clustered environment or to add features to an existing SQL Server 2008 instance.
New SQL Server failover cluster installation
Launch a wizard to install a single-node SQL Server 2008 failover cluster.
Add node to a SQL Server failover cluster
Launch a wizard to add a node to an existing SQL Server 2008 failover cluster.
Upgrade from SQL Server 2000 or SQL Server 2005
Launch a wizard to upgrade SQL Server 2000 or SQL Server 2005 to SQL Server 2008. Before you upgrade, you should run the Upgrade Advisor to detect potential problems.
Search for product updates
Search Microsoft Update for SQL Server 2008 product updates.
SQL Server 2008

FIGURE 8.24
SQL Cluster Installation screen

On this screen you tell the SQL Server software that you will create a clustered configuration for the SQL Server database back end. This will check that you have all the necessary prerequisites such as the disks, cluster nodes, networks, and clustered DTC resources in place before the actual software is installed. Figure 8.25 shows the configuration of the SQL Server cluster after installing the software on the first node of the Windows failover cluster.

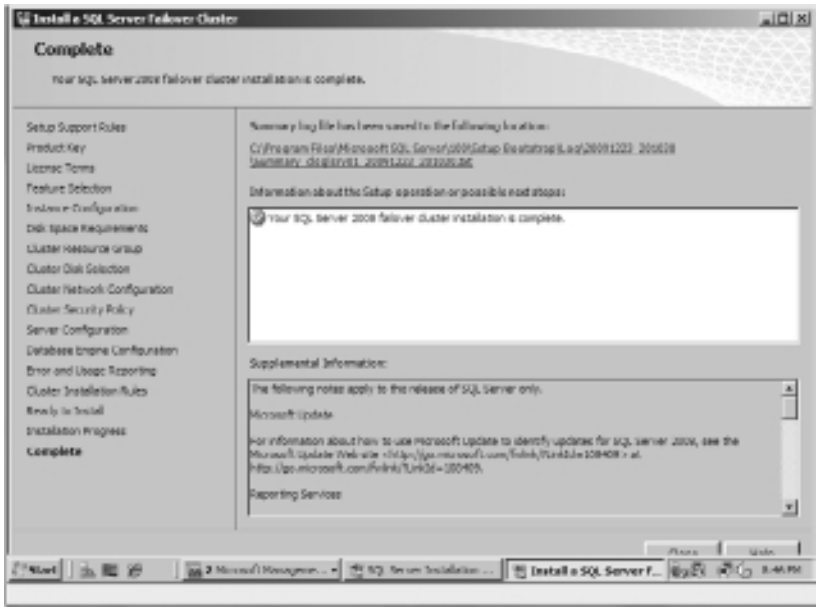
Install a SQL Server failover cluster
Complete
Your SQL Server 2008 failover cluster installation is complete.
Setup Support Rules
Product Key
License Terms
Feature Selection
Instance Configuration
Disk Space Requirements
Cluster Resource Group
Cluster Disk Selection
Cluster Network Configuration
Cluster Security Policy
Server Configuration
Database Engine Configuration
Error and Usage Reporting
Cluster Installation Rules
Ready to Install
Installation Progress
Complete
Summary log file has been saved to the following location:
C:\Program Files\Microsoft SQL Server\100\Setup Bootstrap\Log\20091223 201628\Summary_dbgserver1_20091223_201628.txt
Information about the Setup operation or possible next steps:
Your SQL Server 2008 failover cluster installation is complete.
Supplemental Information:
The following notes apply to the release of SQL Server only.
Microsoft Update
For information about how to use Microsoft Update to identify updates for SQL Server 2008, see the Microsoft Update Web site <http://go.microsoft.com/fwlink/?LinkId=108409>.
Reporting Services
Close
Help
start
Microsoft Management...
SQL Server Installation...
Install a SQL Server F...
8:46 PM

FIGURE 8.25
Successful completion of the first node

At this point, the SQL Server cluster is formed with just the first node of the Windows failover cluster. A single-node cluster is pretty much useless because the failure of that node would constitute a total failure of the SQL Server failover cluster. At this point, you should install any service pack and critical updates on the first node of the SQL Server cluster. Figure 8.26 shows the addition of the second node of the Windows failover cluster to the SQL Server failover cluster configuration.

File View VM
SQL Server Installation Center
Planning
Installation
Maintenance
Tools
Resources
Advanced
Options
New SQL server stand-alone installation or add features to an existing installation
New SQL server failover cluster installation
Launch a wizard to install a single-node SQL Server 2008 failover cluster.
Add node to a SQL Server failover cluster
Upgrade from SQL Server 2000 or SQL Server 2005
Search for product updates
Search Microsoft Update for SQL Server 2008 product updates.
Start
Server Manager
SQL Server Installatio...

FIGURE 8.26
Adding the second Windows failover cluster node as a SQL Server failover cluster node

Adding the second node to the SQL Server failover cluster configuration is pretty much the same as creating the cluster. All of the same checks are performed, and the installation will proceed as it did during the first node of the SQL Server failover cluster install. Just before the installation is completed, the SQL Server software checks that both nodes of the SQL Server failover cluster are running the same versions and patch levels. Because you installed the service pack on the first node of the SQL Server failover cluster, you will have to install the service packs and critical updates onto the second node of the SQL Server failover cluster after the installation is complete.

Once the installation has successfully completed, you will have a two-node SQL Server failover cluster that you will use as the back-end database for the SharePoint Server farm. Figure 8.27 shows the successful completion of the second Windows failover cluster node to the SQL Server failover cluster.

With the SQL Server failover cluster successfully completed, you should test the functionality of the cluster to ensure that all resources and services will properly fail over. Figure 8.28 shows the configuration and the owner of the resources of the SQL Server failover cluster.

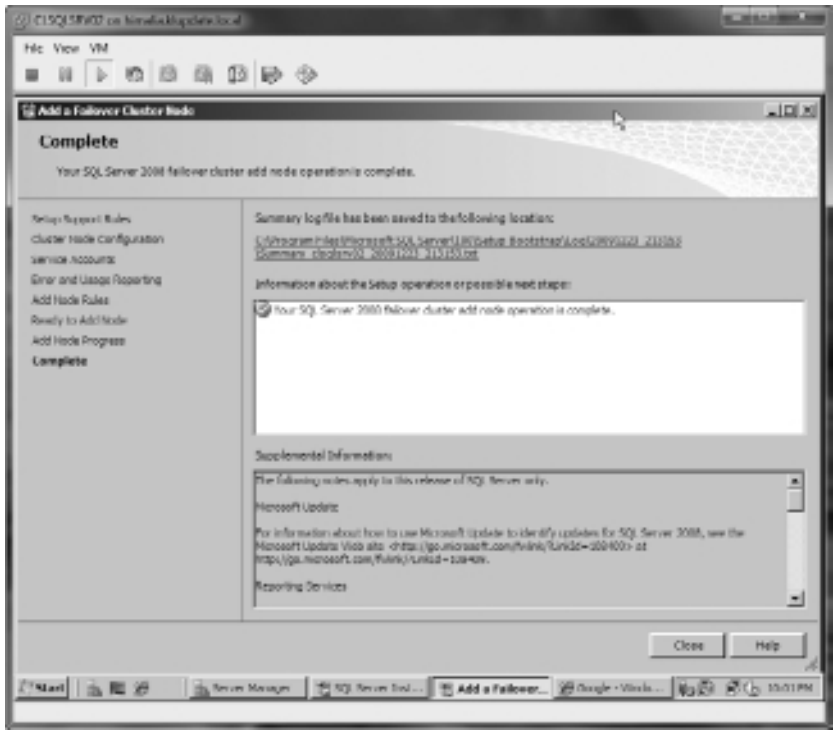
File View VM
Add a Failover Cluster Node
Complete
Your SQL Server 2008 failover cluster add node operation is complete.
Setup Support Rules
Cluster Node Configuration
Service Accounts
Error and Usage Reporting
Add Node Rules
Ready to Add Node
Add Node Progress
Complete
Summary log file has been saved to the following location:
Information about the Setup operation or possible next steps:
Your SQL Server 2008 failover cluster add node operation is complete.
Supplemental Information:
Microsoft Update
Reporting Services
Close
Help

FIGURE 8.27
Successful addition of the second node

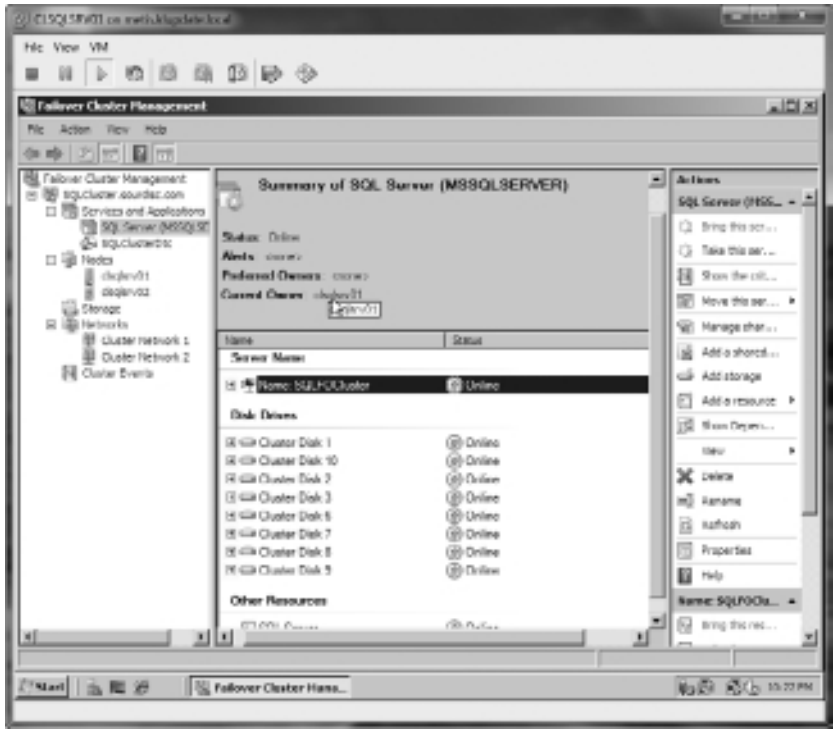

FIGURE 8.28
SQL Server failover cluster before failover

In Figure 8.29, you can see that the owner of the resources and services are owned by cluster node CLSQLSRV01. All of the resources and services are online and functioning properly. Also, you can see that there are no warning events or errors visible. Any errors or warnings that the cluster is experiencing will appear as red Xs or yellow exclamation points at the top of the cluster management utility.

FIGURE 8.29
SQL Server failover cluster after failover

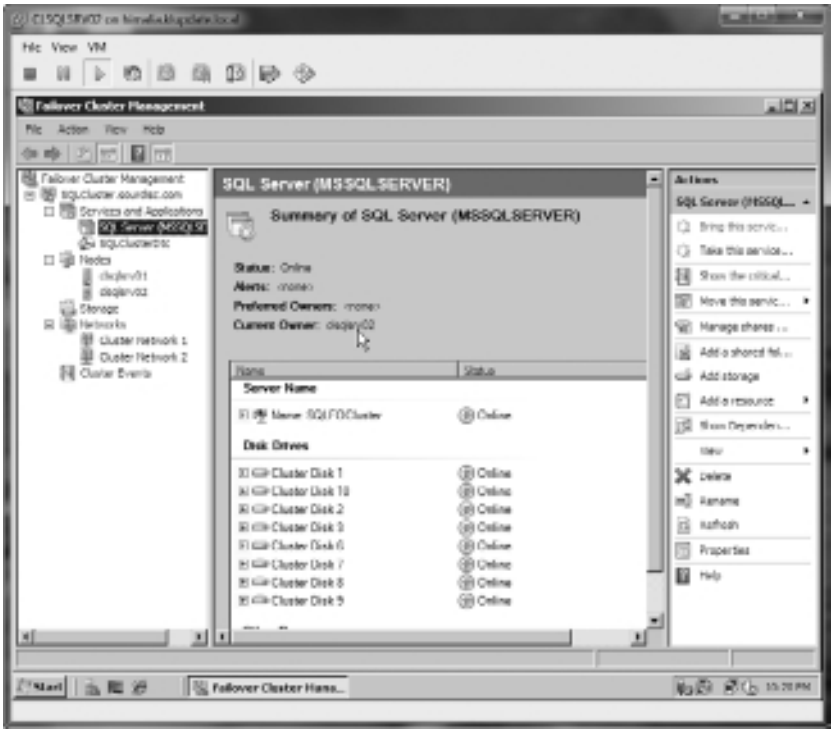

You can see in Figure 8.29 that the SQL Server services and its resources have been moved to the second node of the SQL Server failover cluster. Also, all services and resources are online and functioning properly. You can now start to configure other roles in the SharePoint Server farm because the SQL Server failover cluster database back end is functioning properly and ready to receive the farm's databases.

Configuring the SharePoint Central Site Administration Virtual Machine

The SharePoint Central Site Administration utility is where you administer and perform maintenance on the SharePoint farm. It is here where you will add new servers to the farm when needed, configure additional applications, and configure additional add-in document filters. Using the guidance from Chapter 3, you can build and configure the virtual machine that will host the SharePoint Server Central Site Administration utility as well as the Windows Network

Load Balancing Administration utility. The virtual machine consists of the following hardware (Figure 8.30):

- Two vCPUs.
- Four GB of virtual machine memory. We started with 4 GB of memory because this server role is hosting the SharePoint Central Site Administration as well as the Windows Network Load Balancing utility.
- Two virtual network interface cards: one public network interface card for client access and cluster management and one heartbeat network interface card for cluster communications. As shown in Figure 8.29, there are two separate virtual switches for the cluster configuration.
- Two virtual disks: one for the operating systems and one for the SharePoint Server binaries.

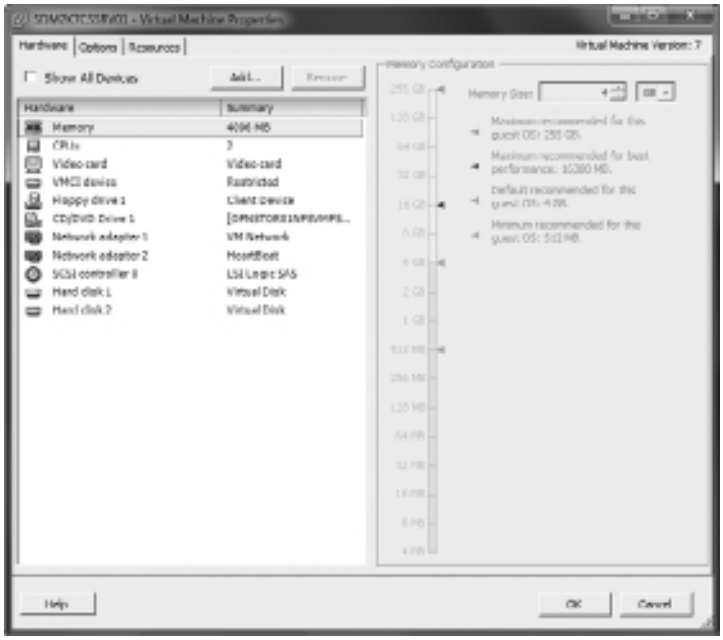

FIGURE 8.30
Virtual machine configuration

In this example, you have not added any additional or special hardware to the virtual machine. After the virtual machine is configured, you installed Windows Server 2008 Enterprise edition and all service packs and critical updates on the operating system. You then installed the web server role and added the Windows Network Load Balancing and .NET 2.0, 3.0, and 3.5 features to the virtual machine. The time required to create this virtual machine configuration can be greatly reduced through the use of the virtual machine template guidance presented in Chapter 3 and adding the following server roles and features to the template:

- Web server role
- Windows Network Load Balancing feature
- .NET features (v2.0 to v3.5)

You could also add the SharePoint binaries in a directory created on the operating system drive. This will eliminate the need to mount a CD or ISO image to install the SharePoint Server software (Figure 8.31).

FIGURE 8.31 SharePoint required server roles and features

This figure shows all the server roles and features that should be installed before attempting to install the SharePoint Server binaries onto the virtual machine. All the required service packs and critical updates should be installed and the virtual machine restarted before you attempt to install any SharePoint Server software.

Configuring the SharePoint Index Server Role

The SharePoint index server role can be a tricky role to virtualize because it is the server role that is most responsible for indexing the SharePoint content. This can be a resource-intensive role in that it requires and uses CPU, memory, and disk resources. It is therefore important that you give this role the required resources as the database back-end server role. Failure to give this role the required resources will result in the performance degradation of the farm. In addition, search query performance will suffer, and user content results will be out-of-date or unavailable altogether. Because this is a proof of concept for Source Discovery and will not initially host all of the user content, we will show how to build the virtual machine using the guidance from Chapter 3 but with the same software additions that we used for the SharePoint Central Site Administration virtual machine minus the Windows Network Load Balancing software. You can start with the following resources (Figure 8.32):

- Two vCPUs
- Four GB of virtual machine memory
- One virtual network interface card (vNIC) for client and server access
- Two virtual disks (one for the operating system and one for the index corpus); the corpus disk is sized for up to 25 percent of the content disk

You can start with two vCPUs and 4 GB of memory for the index server role, and as the corpus grows and the demand for more resources grows, you can adjust the virtual CPU and

memory requirements with just a reboot. With the vSphere platform, you will be able to allocate up to eight virtual CPUs and up to 255 GB of virtual machine memory if needed without provisioning a new server, and you will also gain all the benefits associated with virtualizing this server role.

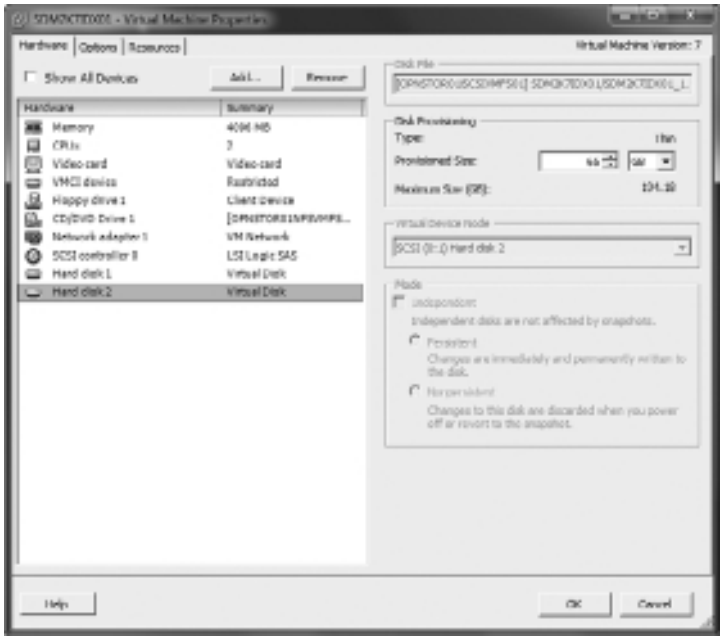

FIGURE 8.32
Index server role virtual machine configuration

Although there can be an index server role for the farm only, it is still a great idea to maintain a template for this server so that provisioning a new index server role virtual machine can take place in a matter of minutes vs. hours or days. You can then think of the index server role virtual machine as a sort of disposable virtual machine.

Configuring the SharePoint Web Front-End Server Role

The SharePoint front-end web server for the proof of concept is built using the same guidance in Chapter 3; it is a mirror virtual machine of the SharePoint Central Site Administration virtual machine. You can use the same template used for the Central Site Administration virtual machine with all the roles and features, including the Windows Network Load Balancing feature. Figure 8.30 shows the Central Site Administration virtual machine configuration that you can use for the front-end web server virtual machines.

- Two virtual CPUs (vCPUs).
- Four GB of virtual machine memory. We started with 4 GB of memory because this server role is hosting the SharePoint Central Site Administration as well as the Network Load Balancing utility.

- Two virtual network interface cards: one public network interface card for client access and cluster management and one heartbeat network interface card for cluster communications.
- Two virtual disks: one virtual disk is provisioned for the operating systems, and one virtual disk is provisioned for the SharePoint Server binaries.

There is not too much difference in the configuration used to build the virtual machine hosting the Central Site Administration utility and the one for the WNLB configuration. Figure 8.33 shows the virtual machine configuration for the SharePoint front-end web server role. This configuration is duplicated to provide load balancing, redundancy, and availability for the front-end web server role.

FIGURE 8.33
SharePoint front-end web server role

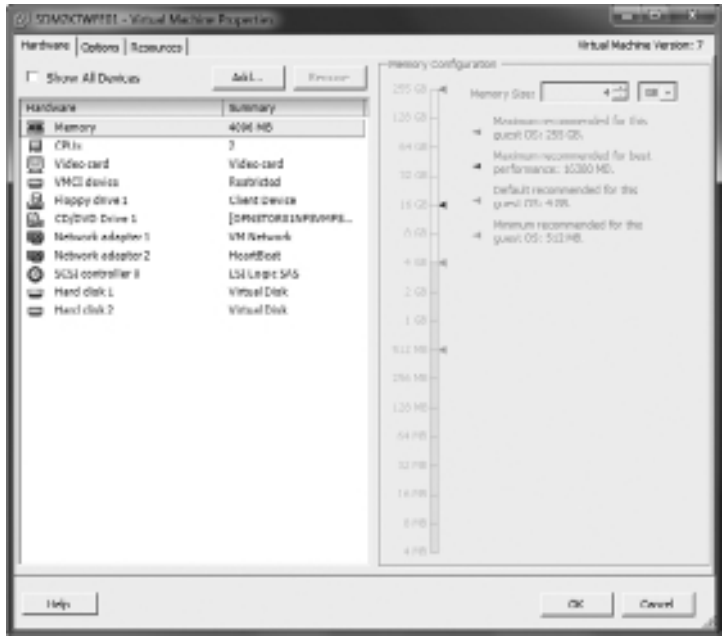

With all of the SharePoint Server role virtual machines configured, it is now time to install the software.

Installing the SharePoint Software

Installing the SharePoint is a two-part process:

1. Install the software onto all of the virtual machines except the SQL Server cluster.
2. If you have included the software binaries in the virtual machine template, then you will run the installer from the directory within the virtual machine.

Here we're mounting an ISO image that is an exact copy of the Microsoft SharePoint Server DVD, so follow these steps:

1. From the SharePoint Central Site Administration utility's virtual machine, run the installer. You are presented with the product entry screen.
2. This is where you enter the product key required to continue with the installation. SharePoint is licensed by the server, so each individual virtual machine running the SharePoint services and features will require a license or product key (Figure 8.34).

FIGURE 8.34
SharePoint product key screen

3. After you have entered the product key in the field, click the Continue button to continue with the installation.
4. On the next screen, agree to the end user licensing agreement (EULA), and then click the Continue button.
5. On the next screen in the installation, you are given a choice of whether you want to perform a basic or advanced installation. Because a SharePoint server role is defined by the features and services running on a server, click Advanced.

We recommend that you choose an Advanced installation for each server because it gives you the flexibility of choosing the SQL Server on which the databases will be created.

6. After you have chosen the Advanced option, you are then presented with the Server Type screen. You can choose from three options (Figure 8.35):
   - Complete: Installs all components. (This provides the most flexibility.)
   - Web Front End: Installs components for the web front end (a single role installation).
   - Stand-alone: Installs all components on a single server (a simple installation). This option does not allow the creation of a multiserver farm.

TIP We recommend performing a full installation of all the components on all virtual machines because this provides the most flexibility in what services run on what virtual machines. This will also reduce the need to reinstall or modify a server role within the SharePoint Server farm.

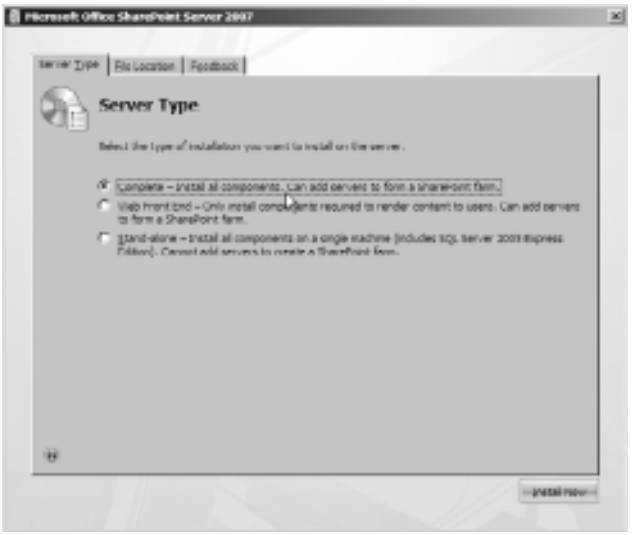
Microsoft Office SharePoint Server 2007
Server Type
File Location
Feedback
Server Type
Complete – Install all components. Can add servers to form a server farm.
Web front end – Only install components required to render content to users. Can add servers to form a SharePoint farm.
Stand-alone – Install all components on a single machine (includes SQL Server 2005 Express Edition). Cannot add servers to create a SharePoint farm.
Install Now

FIGURE 8.35
Installation configuration options

7. After you select the Complete installation type, select the File Location tab. On this tab, you select where to place the SharePoint binaries required to run the SharePoint services and features in the farm. Figure 8.36 shows where we placed the binaries in the operating system virtual disk and the index files on the second provisioned virtual disk in the virtual machine. Configuring the virtual machine in this manner will eliminate contention between the operating system operations and the indexing operations performed by the SharePoint application.

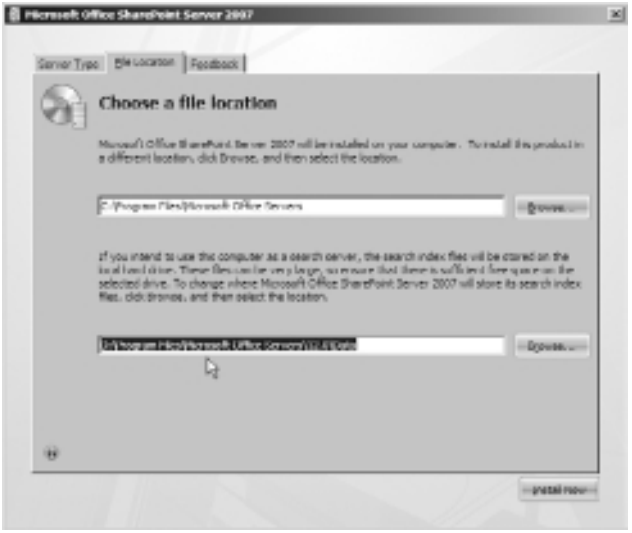

FIGURE 8.36
SharePoint file locations

After you have instructed SharePoint where to place the files for the SharePoint binaries and indexing files (if the server role is the index role), you install the SharePoint application onto the virtual machine. After the installation has completed, you are then presented with the SharePoint configuration wizard screen (Figure 8.37).

FIGURE 8.37
Successful SharePoint installation

At this point, you can choose to install the SharePoint binaries onto the remaining five servers that will make up the SharePoint farm. Repeat the same operation to accomplish the installation of SharePoint throughout the farm.

Configuring Windows Network Load Balancing

Using Windows Network Load Balancing is a good alternative to using a hardware load-balancing solution. For smaller SharePoint farms and development and test environments, this can provide excellent cost-savings benefits. The downside to using WNLB is that the load balancing is handled by the networking stack, which in large SharePoint farms can introduce some performance degradation in the front-end web servers.

TIP For larger SharePoint farm deployments where there are more than two web front-end server roles, we recommend using a hardware load-balancing solution because all of the load-balancing operations will occur on the hardware and not the virtual machine.

Using the configuration of the SharePoint Central Administration/WNLB and web front-end virtual machines, we will show how to create a load-balancing configuration that utilizes the second provisioned virtual NIC in the virtual machine:

1. From the Start menu, select Administration Tools ➢ Network Load Balancing Manager to open the load-balancing utility.
2. Under the Cluster ➢ New option, create a new load balancing cluster.

3. In the Host field, enter the name of the host, and click the Connect button to connect to and configure the first host in the Windows load-balancing cluster. In Figure 8.38 we have entered the first web front-end host.

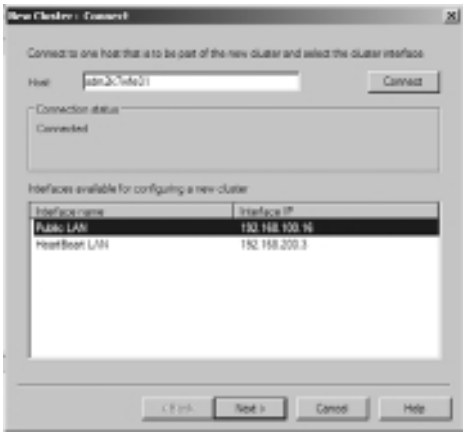

FIGURE 8.38
Entering the first web front end server to create the WNLB cluster

You can see that after we have connected, the second virtual NIC is detected and configured automatically as the WNLB NIC that will be used for the load-balancing operations. The IP address for the second virtual NIC was entered ahead of time after the configuration of the Windows operating system.

WARNING When configuring the NLB virtual NIC, do not select the Network Load Balancing (NLB) check box because doing so will prevent the Windows Load Balancing Manager from being able to access the virtual NIC for configuration.

4. Click the Next button to continue the installation.

5. Enter the IP address and subnet mask for the cluster. Figure 8.39 shows the IP and subnet addresses that we have chosen for the load balancing cluster.

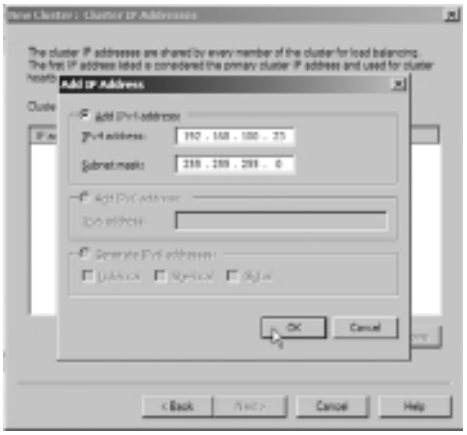

FIGURE 8.39
Choosing the cluster IP address

6. Clicking the OK button takes you to the next screen where you must enter a fully qualified DNS name for your WNLB cluster.
7. After entering the fully qualified DNS name for the cluster, select the Unicast option for the cluster operation mode. This will enable the WNLB cluster to utilize the second provisioned virtual NIC.
8. On the next screen, you have the option of specifying which traffic will go through the WNLB cluster by specifying the different ports for the different SharePoint services. We have left these at the default (see Figure 8.40). You can continue by clicking the Finish button.

FIGURE 8.40
Forming the Windows NLB cluster

At this point, you have created and formed the Windows NLB cluster with a single SharePoint web front-end server. In this example configuration, you have a cluster but a single point of failure and no load balancing or redundancy for the SharePoint web front-end server role. If you were to experience a web front-end server failure, the entire SharePoint farm would be inaccessible to the Source Discovery end users. This would not be an acceptable solution and would defeat the purpose of deploying WNLB in the first place. So, to continue, you can add the second web front-end server to the existing WNLB cluster configuration by selecting Cluster ➢ Add Host To Cluster. Figure 8.41 shows the connection to the second SharePoint web front-end server.

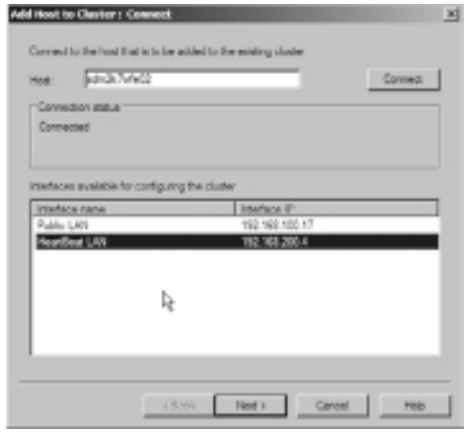

FIGURE 8.41
Adding the second SharePoint web front end server to the Windows NLB cluster

As you can see, the process is the same as adding the first SharePoint web front-end server to form the cluster. After adding the second host, the cluster will go through a cluster transition state known as *convergence* in which both cluster nodes are synchronized with one another. This could take up to a minute or more to complete convergence. Once complete, all the nodes of the WNLB cluster should be fully operational and display a green status. Figure 8.42 shows the Windows NLB cluster state after convergence has completed.

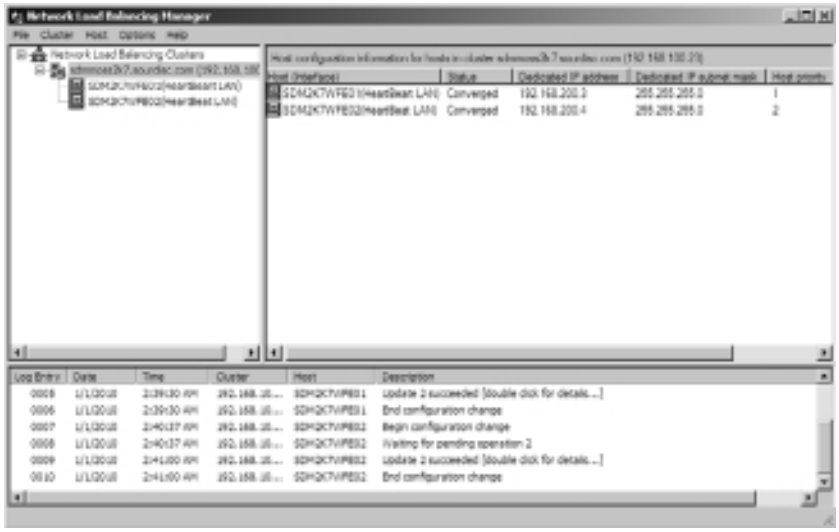

FIGURE 8.42
A fully operation Windows NLB cluster

After configuring the SharePoint farm, you can test the functionality of the Windows NLB cluster. To do this, power off the first SharePoint web front end and connect to the Windows NLB cluster name to ensure that you can access Source Discovery's SharePoint farm websites. Once your tests are completed, you're finished with the Windows Network Load Balancing cluster configuration.

LOOKING FORWARD

Any discussion about SharePoint 2007 would not be complete without referencing the much-anticipated next major release of SharePoint 2010. As of this writing, Microsoft has announced that the planned release of SharePoint 2010 will be in the first half of 2010 and has made the SharePoint 2010 beta available to the public. Let's look into some of the key features of SharePoint 2010 and their impact on virtualization from an information technology professional perspective based on the current SharePoint 2010 beta features:

- The minimum operating system requirement for SharePoint 2010 is Windows Server 2008 SP2 (64-bit) or higher. If you are still running on 32-bit Windows Server or on Windows Server 2003, it is time to start looking for an operating system upgrade.
- The minimum SharePoint 2007 service pack level for upgrade to SharePoint 2010 is SP2. If you are still running on SharePoint 2007 SP1, it is time to upgrade your service pack.

- The Shared Services Model in SharePoint 2007, which had several limitations including SSP association overhead for all web applications, is now rearchitected with a flexible service application model. This rearchitecture has resulted in several new SharePoint built-in databases and server roles.
- SharePoint 2010 has better health and monitoring tools. The Best Practices Analyzer Tool in SharePoint 2007 is now integrated with the product stack under the Health Analyzer. The Developer Dashboard feature (though the name is misleading) can provide an end-to-end insight of each web page starting from web server to database transaction. Unified Logging Server and Performance Control/Throttling features are some of the additional features for IT Pros to better manage SharePoint 2010.
- The biggest news on the High Availability front is, SharePoint 2010 built-in support for Mirror failover. You don't have to write those complex scripts, which is done today for SharePoint 2007 to re-act to database mirror failover.
- Other features like PowerShell 2.0 support, LoadTest Kit 2010, Business Connectivity Services, Multi-tenant hosting, support for large lists, external BLOB storage are also worth mentioning.

With SharePoint 2010 support for scaling out servers including index crawl, you will possibly need more servers to maintain the minimum high-availability requirements of a SharePoint 2010 farm. Businesses can expect to implement all of the new features of SharePoint 2010 on their existing physical infrastructures. Expect 15 to 20 databases minimum. Virtualization can play a great role in consolidation and can provide a smoother upgrade at a minimum to SharePoint 2010.

Summary

In this chapter, we covered SharePoint governance and why you would need governance. SharePoint governance is necessary to ensure that the SharePoint solution that you design and implement will meet your organization's changing needs. Governance will also ensure that the content data that is critical to your organizational operations is handled properly and that only the individuals who require access to content have access to the appropriate content. Finally, governance will help you design and implement a SharePoint Server solution that will be utilized effectively.

We also covered SharePoint Server roles and their responsibilities as well as the characterization and the suitability for virtualizing the various SharePoint Server roles. We then moved on to talk about sizing a SharePoint Server farm by using the Microsoft System Center Capacity Planner with a SharePoint plug-in. In this sizing tool, the solution is designed about generic server and storage hardware specifications.

We then showed how to design a SharePoint Server 2007 farm solution using the HP SharePoint sizing tool. In the HP SharePoint 2007 sizing tool, the solutions are designed around HP server and storage hardware solutions. The HP solution also has an update function to synchronize current pricing information to allow you to price your SharePoint 2007 solution.

Finally, you followed the fictitious company Source Discovery along its journey of total Microsoft data-center virtualization. We showed how to deploy a proof-of-concept SharePoint 2007

farm consisting of six virtual machines. You learned how to implement high availability for the SharePoint Server 2007 web front-end server role using Windows Network Load Balancing and Microsoft Failover Clustering for the database back-end server role. In the proof of concept, we discussed virtual machine configuration for each server role, and we showed how to install the SharePoint binaries for all virtual machines in the proposed farm.

After installation and configuration, some of tasks completed by the consultant were:

- Moving SharePoint logs and index to a different drive using SharePoint Server stsadm tool
- Running the SharePoint Best Practice Analyzer tool to validate the configuration
- Testing their SharePoint Server 2007 farm configuration with their proposed workload requirements
- Scripting the database and log file movements from the default disks to the clustered disks

Source Discovery populated the farm with a 100 GB portion of its 200 GB content data. The solution is being utilized alongside the existing collaboration/document management solution. After careful evaluation of the virtualized implementation, Source Discovery will transition the proof of concept into production by populated the remaining content data into the virtualized solution.

Many VMware customers have successfully virtualized SharePoint Server 2007 using the vSphere platform. All server roles of the SharePoint Server farm scale very well and through virtualization can provide considerable benefits in cost savings, flexibility, and efficiency. Because Microsoft SharePoint Server 2007 provides scalability options for the web server front-end server role and the SQL Server database back-end server role, you can virtualize larger farms successfully by paying close attention to information architecture, SharePoint Server boundaries (site topologies, document libraries, and so on), and SQL Server considerations for virtualization.

Chapter 9

Virtualizing Microsoft Remote Desktop Services

Remote Desktop Services enables multiple concurrent users to run several applications on a remote computer. In Windows Server 2008 R2, Microsoft renamed Terminal Services to Remote Desktop Services (RDS) and introduced access to virtual desktops in addition to the traditional session-based remote desktops. RDS encompasses the roles formerly under Terminal Services as well as a new role service to manage and automate access to virtual desktops in a virtual desktop infrastructure (VDI).

In this chapter, we will review the RDS roles and enhancements in Windows Server 2008 R2, key considerations for virtualizing RDS roles, application compatibility checking for RDS deployment, and characterizing and sizing RDS environments. We will also continue with our data center proof of concept by deploying a sample RDS farm for the Source Discovery Corporation. We will focus on the traditional session-based remote desktops provided by RDS in this chapter. The VDI enabled by RDS is beyond the scope of this book.

RDS Virtualization Considerations

Given the multiple server roles for RDS, you must consider workload characteristics for each server individually. Proper RDS role resource planning should be exercised here because any lack of the required resources for a server role will result in an unacceptable end-user experience throughout the RDS farm.

TIP Microsoft provides a detailed infrastructure design and planning guide to architect and implement RDS server farms. You can access this information from the RDS TechNet website at `http://technet.microsoft.com/en-us/windowsserver/ee236407.aspx` or at `http://go.microsoft.com/fwlink/?LinkId=177881`.

In general, RD Gateway, RD Web Access, RD Connection Broker, and RD Licensing are lightweight roles, and you should have no problems virtualizing them. All these RDS roles generally handle short transactions and then pass the more substantial duties to an RD Session Host server.

The RD Session Host role is the most challenging role to virtualize because of its unique workload profile, not only compared to other RDS roles but also relative to other enterprise server applications that are virtualized. Because of the diverse resource requirements for each server role, we will discuss the considerations for virtualizing each RDS role separately in the following sections.

RDS Roles in Windows Server 2008 R2

In Windows Server 2008 R2, all Remote Desktop Services roles have been renamed. The following table lists the old and new names for each role.

Previous Name	Name in Windows Server 2008 R2
Terminal Server	Remote Desktop Session Host (RD Session Host)
Terminal Services Licensing (TS Licensing)	Remote Desktop Licensing (RD Licensing)
Terminal Services Gateway (TS Gateway)	Remote Desktop Gateway (RD Gateway)
Terminal Services Session Broker (TS Session Broker)	Remote Desktop Connection Broker (RD Connection Broker)
Terminal Services Web Access (TS Web Access)	Remote Desktop Web Access (RD Web Access)

RDS expands the role of RD Connection Broker to provide user access to the traditional session-based remote desktops and the newly introduced VDI. RD Connection Broker complements shared RDS infrastructure components in Windows Server 2008, such as RD Web Access or RD Gateway. Windows Server 2008 R2 also introduces a series of platform enhancements for remote desktop users, such as support for multiple physical monitors; redirection of multimedia and 3D content, including Vista Aero; and enhanced, bidirectional audio support.

RDS also introduces a new service role, not listed in the previous table, called Remote Desktop Virtualization Host (RDVH), which is based on Microsoft Hyper-V technology to enable a VDI solution. RDVH integrates with the Hyper-V virtualization platform to provide access to virtual desktops running as virtual machines. Users can be assigned a unique personal virtual desktop or redirected to a shared VM pool from which a virtual desktop is dynamically assigned. RDVH leverages RemoteApp and Desktop Connection and RD Web Access to publish these virtual desktop resources to users. It uses an RD Session Host as a "redirector" together with RD Connection Broker for serving these VM-based requests.

A remote desktop (RD) client gets connection information from the RD Web Access server in an RDS solution. If an RD client is outside a corporate network, the client connects through an RD Gateway. If an RD client is internal, the client can then directly connect to an intended RD Session Host or RD Virtualization Host once RD Connection Broker provides the connection information. In both cases, RD Connection Broker plays a central role to make sure a client gets connected to the correct resource. With certificates, an RDS farm can configure digital signing and single sign-on to provide a great end-user experience with high security. The following illustration shows the various service roles in an RDS farm.

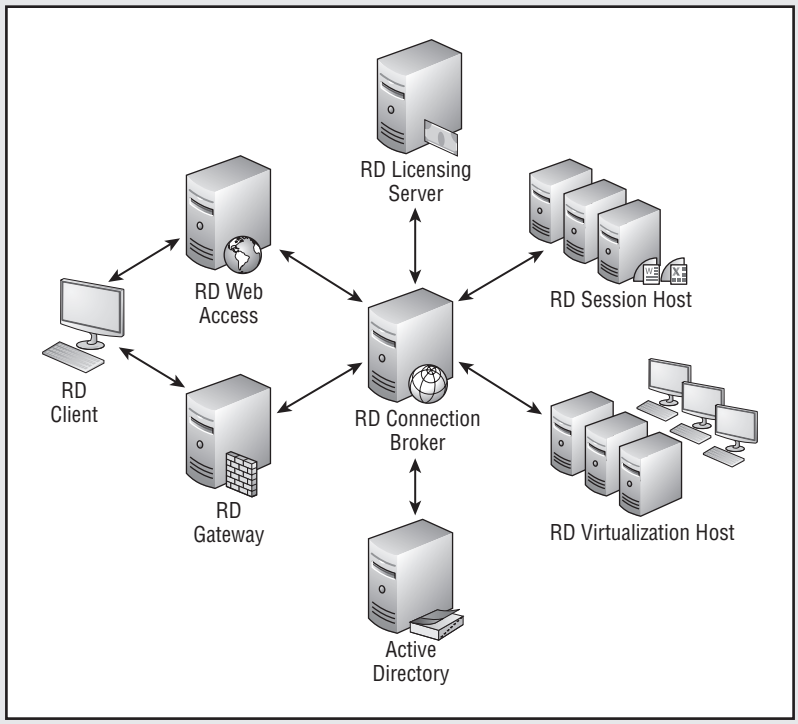

Deploying VDI using RDVH is a really a desktop virtualization topic and can take up an entire book. Therefore, we will not cover VM-based RDS in this chapter and will focus only on the traditional session-based RDS served by RD Session Host.

RD Session Host Role

The RD Session Host provides the functionality that was known as Terminal Services in previous releases. Using the RD Connection Broker, RD Session Host servers provide access to applications—either by providing access to the server's desktop or by using RemoteApp. RemoteApp, introduced in Windows Server 2008, allows you to publish specific programs to remote users, rather than enable access to entire desktops. Remote users can run these RemoteApp programs alongside their local applications using the same look and feel as their local desktop. A user accessing multiple RemoteApp programs running on same RD Session Host share a single RDP connection to the host. You can make RemoteApp programs available to remote users either through RD Web Access interface or by distributing `.rdp` files or Windows Installer packages for these programs. When RemoteApp is used with Windows 7 clients, you can also use the RemoteApp and Desktop Connections feature to distribute RemoteApp program links.

Multiple RD Session Host servers can be deployed in an RDS farm to provide aggregate bandwidth, load balancing, and redundancy in the event of RD Session Host server failures. This requires every RD Session Host in a load-balancing group to be identically configured with the same applications.

The RD Session Host (RDSH) role is the most challenging role to virtualize among all RDS roles. The RDSH server exhibits unique workload characteristics compared to other server applications such as web, database, and email servers. An RD Session Host has to efficiently juggle between multiple users and all the applications each user runs. An RD Session Host server must do the following:

- Ensure that new users experience acceptable response times as they log on while continuing to service currently connected users
- Run many instances of same applications as efficiently as possible
- Evenly distribute CPU resources across all sessions so that no one session is dominating the available server resources
- Keep track of physical memory available and use it as efficiently as possible
- Effectively manage any storage I/O associated with the page file

Given these requirements, we will discuss the considerations in virtualizing the RD Session Host in the following ways:

- Windows architecture
- Processor architecture
- Application diversity
- Usage profile

Windows Architecture

One of the major issues when running an RD Session Host server on the 32-bit Windows architecture is the limited space available for the operating system kernel virtual address space. The 32-bit operating system reserves a 2 GB virtual address space for the kernel data structures. This virtual address space is shared by all processes that are running on the system. There are three areas in the kernel address space that have a significant impact on RD Session Host server scalability: the paged pool area, the system page table entries (PTEs) area, and the system cache area. All these three areas have lower hard limits in a 32-bit Windows environment. When this kernel address space is exhausted, no new process (or any system object) can be created. This means that new users can no longer log on to the system and that the currently logged-on users are severely impacted in terms of performance. The kernel virtual address space size on x86-based systems limits the number of sessions supported regardless of CPU, memory, and I/O capacity of the RD Session Host server machine.

In comparison, the 64-bit architecture provides a significantly larger virtual address space for the kernel data structures (8 TB). This increased kernel virtual address space hosts a 128 GB paged pool and 128 GB system page table entries (PTEs). Table 9.1 compares the relevant kernel virtual address space limits in 32-bit and 64-bit Windows Server 2003.

Use Virtualization to Scale 32-bit RDS Deployments

When user scalability is impacted because of the kernel memory limits in 32-bit Windows OS, you should leverage virtualization to scale your 32-bit RDS deployments. By running multiple RD Session Host virtual machines on a vSphere host, you can scale to a much larger number of users compared to physical RDS deployment. Although each RD Session Host virtual machine will run slightly fewer users because of some virtualization overhead, multiple RD Session Host virtual machines together will provide a higher cumulative user count.

Table 9.1: Comparison of x86 and x64 Memory Areas

Architectural Component	x86-Based Windows Server 2003 (Terminal Server Optimized)	x64-Based Windows Server 2003
Virtual memory	4 GB	16 TB
Kernel virtual address space	2 GB	8 TB
Paged pool	260–480 MB	128 GB
Nonpaged pool	256 MB	128 GB
System cache	1 GB	1 TB
System PTEs	About 900 MB	128 GB

Thus, 64-bit architecture essentially removes kernel virtual address space limitations that affect the number of sessions that are supported by the operating system in the 32-bit architecture. Although the actual user session limit is affected by the number of applications running in the session, by the resource consumption of each application, and by the user activity patterns, the 64-bit RD Session Host will scale much better when compared to a 32-bit one.

Windows Server 2008 makes further improvements to the three kernel memory areas. The earlier versions of the Windows operating system had predefined size limits for these areas. Windows Server 2008 makes these sizes flexible and can dynamically adjust these limits as needed.

WARNING Do not use the /3GB switch on RD Session Host. If you use the /3GB switch to allocate more address space to user mode processes, the kernel memory space is limited to 1 GB, and you decrease the number of available system PTEs. This will result in kernel memory exhaustion sooner, causing the RD Session server to become unresponsive and thus limiting user scalability for RD Session Host.

Processor Architecture

RD Session Host is a multiuser environment, with each user session running a number of processes. Consequently, this environment exhibits a high number of context switches between various processes, lots of process creations and destructions, and high paging activity. All of these make the RD Session Host a memory management unit (MMU)–intensive workload. Frequent page table updates associated with process context switching need to be maintained per processor, so using servers with multiple processors puts additional load on the MMU.

Memory management in virtual machines is challenging compared to physical machines, especially when it comes to virtual memory address translation. In physical machines, the operating system uses page tables to translate the memory address from the application's "virtual" space into the machines' physical memory addresses. Similarly, in virtual machines, guest virtual memory addresses are translated to guest physical addresses using the guest OS's page tables. However, because the guest OS does not have access to the physical machine memory, vSphere is responsible for the final translation to machine physical memory addresses. The binary translation virtualization technique traditionally used by ESX maintains a copy of guest OS page tables called shadow page tables for each virtual machine. Creating and maintaining the shadow page tables requires both CPU and memory resources. This can result in significant performance overhead on RD Session Host, which exhibits high page table activity.

Both AMD and Intel have introduced hardware-assisted memory management capabilities to alleviate this situation. Hardware-assisted memory management capabilities from AMD and Intel are called Rapid Virtualization Indexing (RVI) and Extended Page Tables (EPT), respectively. This support consists of a second level of page tables implemented in hardware. These page tables contain guest physical to machine memory address translations. ESX 4.0 introduces support for the Intel Xeon processors that support EPT. Support for AMD RVI has existed since ESX 3.5. This hardware support eliminates the need for maintaining shadow page tables and the associated performance overhead. Consequently, RD Session Host virtual machines using hardware MMU virtualization are able to scale to significantly higher user sessions. Using such processors is a must when virtualizing RD Session Host.

NOTE VMware has published performance studies showing that AMD RVI and Intel EPT can improve user count on an average of 30 percent in Remote Desktop Services environments. See `www.vmware.com/pdf/Perf_ESX_Intel-EPT-eval.pdf` and `www.vmware.com/pdf/RVI_performance.pdf`.

vSphere will automatically set hardware MMU virtualization if available; you can also explicitly set it using VM properties (Figure 9.1).

You can also check the `vmware.log` file for a virtual machine to make sure that RD Session Host VM is using hardware MMU virtualization. The following code shows the section of `vmware.log` displaying the virtualization technique used by the virtual machine:

```
Feb 18 22:03:09.512: vmx| MONITOR MODE: allowed modes          : BT HV HWMMU
Feb 18 22:03:09.513: vmx| MONITOR MODE: user requested modes   : BT HV HWMMU
Feb 18 22:03:09.513: vmx| MONITOR MODE: guestOS preferred modes: HWMMU HV BT
Feb 18 22:03:09.513: vmx| MONITOR MODE: filtered list          : HWMMU HV BT
Feb 18 22:03:09.514: vmx| HV Settings: virtual exec = 'hardware'; virtual mmu =
'hardware'
```

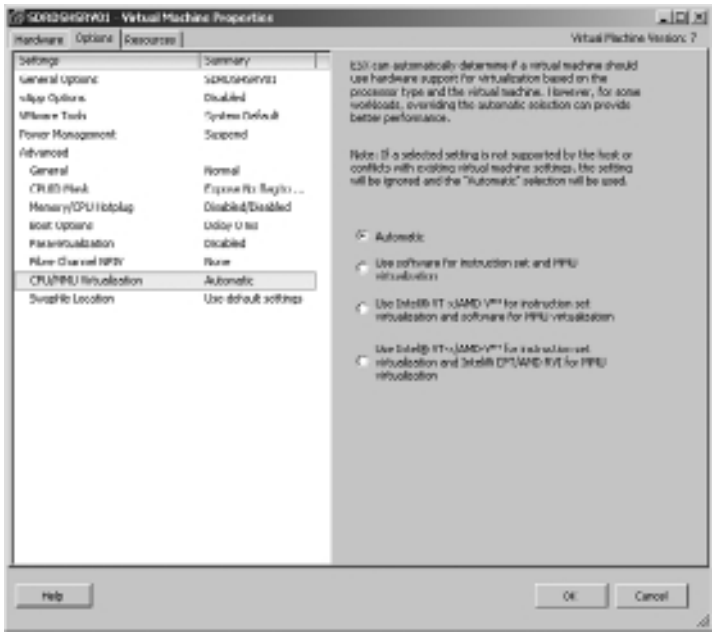

FIGURE 9.1
Setting Hardware MMU virtualization

TIP VMware has published detailed white papers on vSphere Virtual Machine Monitor mode at `www.vmware.com/files/pdf/software_hardware_tech_x86_virt.pdf` and `www.vmware.com/files/pdf/perf-vsphere-monitor_modes.pdf` that provide more details on this topic.

Application Diversity

Another key characteristic that makes the RD Session Host role tricky to virtualize is the wide variety of applications that it runs. We often describe an RD Session Host environment as a custom-made pizza. You can take an RD Session Host server (the pizza base) and then add your applications (the pizza toppings) to create your own pizza. As a result, every RD Session Host deployment is unique. The applications that customers run on RD Session Host range from off-the-shelf Microsoft Office applications to home-grown, custom applications that bring their own set of performance requirements and challenges. The workload profile of applications running will greatly determine the performance of a virtualized RD Session Host.

With the introduction of Windows Server 2008 R2, the Windows Server operating system is compatible only with 64-bit processor architectures. Although this helps the scalability of RD Session Host, as discussed in the previous section, it can be a different story for your applications. Migrating your applications to a 64-bit architecture is not always easy and may not necessarily benefit your applications; 64-bit Windows environments do not support mixed 64-bit/32-bit processes, 32-bit kernel code, and any DOS or 16-bit applications. Most 32-bit applications can run using the Windows on Windows 64 (WOW64) emulator but will incur

significantly higher memory usage. You will need to consider these application issues if you plan to deploy a 64-bit RDS environment.

Usage Profile

In addition to the application diversity we discussed in the previous section, the actual usage profile of RD Session Host needs to be carefully considered for successful virtualization. The usage profile comprises the number of users, their usage patterns, and the user locations from where the applications are accessed. These factors can significantly change the end-user experience even for the same application.

Most published performance studies for RD Session Hosts typically use Microsoft Word with predefined user scripts. Some customers tend to set their user scalability expectations based on these published reports. Unfortunately, real-life workloads vary drastically in terms of applications running as well as usage profile. We will discuss this further in a later section, but suffice it to say that a realistic usage profile combined with applications you plan to deploy will determine the performance of a virtualized RD Session Host in your environment.

NOTE It is important that you follow the Microsoft guidelines and recommendations for deploying the RD Session Host server role.

RD Licensing Role

RD Licensing provides the license management system that enables RDS to obtain and manage RDS client access licenses for devices and users connecting to RDS. RD Licensing manages the Remote Desktop Services client access licenses (RDS CALs) that are required for each device or user to connect to a RD Session Host server. You use RD Licensing Manager to install, issue, and track the availability of RDS CALs on a RD Licensing server.

The RD Licensing role makes an excellent candidate for virtualization for several reasons. In fact, RDS administrators have been running this role in VMs for quite some time to make it easier to maintain a backup.

Lightweight service RD Licensing role is a lightweight service that requires very little CPU, memory, or hard disk resources. The RD Licensing server is active only when a client accesses the RD Session Host server for the first time or when the client license expires. Even under a high load, its impact on the performance of the RDS server farm is very low.

Mandatory The RD Licensing role is the most fragile component in the RDS server farm because it is mandatory but cannot be clustered. Also, RDS CALs on an RD Licensing server are tied to that license server.

Server specific In case of failure, you cannot just reinstall RDS CALs on a new license server. For other RDS roles, you can restore a server role to another server by simply installing the role service and restoring the data used by that role service. Because the RDS CALs are tied to the OS-specific information you gave to the Microsoft Clearinghouse upon activation, you cannot do this with RD Licensing server. If you change the RD Licensing server, the certificates handed down to the RD Session Host servers will no longer be valid, and new ones will need to be issued. Virtualizing your RD Licensing server breaks ties to specific hardware.

Because the RD Licensing role service consumes many fewer resources, it can be easily combined with other roles on one VM. We recommend creating a separate VM for this role, however, so you can benefit from vSphere features such as vMotion, DRS, and HA. You can start with a single virtual CPU (vCPU) and 2 GB of memory virtual machine and then increase to two vCPUs if needed. Because installing RDS CALs on an RD Licensing server immediately associates these licenses with that server, it's good to plan your high-availability strategy before you install RDS CALs. For best availability, you should have two license servers in each forest and split your licenses between the two. Furthermore, we recommend you enable vSphere HA on the RD Licensing server VMs to minimize any unplanned downtime.

Virtualizing the RD Licensing server works regardless of whether you use per-user licensing or per-device RDS CALs. Per-user RDS CALs are not assigned but are updated in the user's Active Directory account to reflect the RDS CAL usage. Because RDS CAL allocation information is not stored on the RD Licensing server, any downtime on RD Licensing server is not as critical. When using per-device RDS CALs, a license allocations database is maintained on the RD Licensing server. This can be an issue if the RD Licensing server fails and you have to restore from an older backup. You can alleviate this using vSphere HA, where the only downtime for the RD Licensing server is the time it takes to restart the VM. If needed, you can avoid even that downtime by using vSphere FT.

RD Connection Broker Role

RD Connection Broker expands the Terminal Services Session Broker capabilities included in Windows Server 2008 to provide a unified experience for traditional session-based remote desktops and VM-based virtual desktops. A VM-based remote desktop can be either a personal virtual desktop or part of a shared virtual pool.

RD Connection Broker supports load balancing to manage user load across multiple RD Session Host or RD Virtualization Host servers. It also reconnects a user to an existing session if the person is disconnected. This works for both virtual desktops and RD Session Host server–based sessions.

RD Connection Broker can also be considered a lightweight service because it is used only when a user connects to an RDS farm. RD Connection Broker maintains no state data, and its connection database is temporary. So, it can be successfully virtualized and protected using vSphere HA or FT. We recommend you start with a single vCPU and 2 GB of memory virtual machine and increase to two vCPUs only if needed.

RD Connection Broker provides a load-balancing service for incoming user session requests in an RDS farm and hence needs to be available to avoid any orphan sessions. RD Connection Broker can be clustered using Windows Failover Clustering. However, depending upon your SLAs, vSphere HA can provide you with enough high availability to avoid the complexity associated with clustering.

RD Gateway Role

RD Gateway is an RDS role service that allows authorized remote users to connect to remote desktops or RemoteApp programs on an internal network from any device connected to the Internet. RD Gateway uses the Remote Desktop Protocol (RDP) over Hypertext Transfer Protocol over Secure Sockets Layer (HTTPS) to establish a secure, encrypted connection between remote users on the Internet and the internal network resources. RD Gateway eliminates the need to establish a VPN connection to the corporate network.

RD Gateway is optional and usually placed at the edge of a corporate network. Using a server certificate, RD Gateway offers secure remote access to the RDS infrastructure. It uses two policies in a designated Network Policy Server (NPS) to do so:

Connection Authorization Policy (CAP) CAP is similar to a user authorization list, showing who can access an associated RD Gateway.

Resource Authorization Policy (RAP) RAP is a resource list specifying which devices a CAP user can connect to via an associated RD Gateway.

The RD Gateway role can also be virtualized without any issues. The RD Gateway role provides authentication and encrypted connection for clients accessing an RDS farm from outside a firewall. The key performance requirement is that it needs sufficient bandwidth to handle the number of simultaneous incoming connections. RD Gateway uses few resources, so it can safely be combined with other roles that may be in the perimeter network. As with other lightweight RDS roles, we recommend you start with a single vCPU and 2 GB of memory virtual machine and then increase to two vCPUs only if needed.

The RD Gateway server may maintain some state data in the form of RD CAPs and RD RAPs, but it doesn't have to. You can maintain RD CAPs and RAPs on a central server and have the virtualized RD Gateway servers point to the central location.

RD Gateway servers need to be load balanced using a hardware load balancer or Windows Network Load Balancing (NLB) to ensure that RDS is not inaccessible from the Internet. In either case, you can extend RD Gateway servers' high availability using vSphere HA or FT. Do not deploy all of your RD Gateway virtual machines onto the same vSphere host in order to prevent single points of failure. The use of DRS anti-affinity rules can prevent the RD Gateway virtual machines from migrating to the same vSphere host. Consider cost and performance when deciding whether to use hardware load balancing instead of software load balancing. If you are deploying software load balancing, deploy dedicated virtual network adapters that are connected to two virtual switches (vSwitches), which in turn should be connected to at least two physical network adapters for the public and cluster communications traffic.

RD Web Access Role

RD Web Access, formerly TS Web Access, provides a web interface for users to access remote resources such as session-based remote desktops, session-based remote applications, or VM-based remote desktops through a web browser. By default, RDS publishes resources at `https://RD-WebAccess-server-FQDN/rdweb`, and an RDS administrator only needs to provide this URL information to a user. When a user starts a RemoteApp program, an RDS session is started on the RD Session Host server that hosts the RemoteApp program. If a user connects to a remote desktop, a Remote Desktop Connection is made either to a desktop on an RD Session Host server or to a VM that is running on an RD Virtualization Host server.

The RD Web Access role in an RDS farm can also be easily virtualized. This role is a very lightweight role in that, on average, it barely requires any resources to perform effectively. As with other ancillary RDS roles, we recommend you start with a single virtual vCPU and 2 GB of memory virtual machine and increase to two vCPUs only if needed. You can also extend the RD Web Access server's high availability using vSphere HA or FT. Again, do not deploy all of your RD Web Access virtual machines onto the same vSphere host in order to prevent single points of failure. You can also improve performance by disabling IIS logging if it is not required. IIS logging captures all information about web functions to include errors, warnings, and informational events on the server.

We will cover the creation and configuration of virtual machines for all the RDS roles in more detail in the proof of concept (PoC) of this chapter.

Characterizing Remote Desktop Services Environment

RDS can be deployed in a variety of topologies, with the most common being a single-server deployment in which all server roles are part of a single server or the distributed model (more commonly known as the farm mode). RDS roles are really defined by the services that are running in the server in an RDS farm. There is a lot of flexibility when deploying in the farm mode. You can combine server roles to reduce the number of servers required for your organization. For example, you could combine the Web Access role with the Session Host so that you reduce the number of hops that the information must travel between servers. The applications running on RD Session Host and the user activity further differentiate each RDS environment. It would be fair to say that every RDS deployment is unique.

Whether you are migrating your physical RDS deployment to virtual or have a "virtualize first" policy and are performing a fresh RDS deployment in a virtual environment, we recommend you make a detailed inventory of your RDS deployment. This inventory needs to cover at least the following items:

- RDS role inventory
- Application inventory
- User inventory
- Business assets inventory
- Software inventory
- Hardware inventory

Next we provide an example checklist of inventory items you should track for each of these categories:

Role inventory This gives a big picture of your RDS ecosystem. Knowing what RDS server roles you have and how they communicate can help you track what you need. Listed here are some items to track:

- Number of RD Session Host servers
- Number of RDS farms
- RD Session Host server distribution within these RDS farms
- Number of RD Gateway servers (if more than one, are they load-balanced?)
- Number of RD Connection Broker servers (if more than one, are they load-balanced?)
- Number of RD Web Access servers (if more than one, are they load-balanced?)
- Number of RD License servers available
- RDS licensing model used
- RDS CALs distribution (single or multiple RD License servers?)
- Any RDS server roles combined on the same server

Application inventory This will help you understand the applications that are used in your RDS deployment. You can also use this to track any application compatibility issues as well as licenses you may need. The following are some examples of application inventory items:

- Number of applications
- List of RemoteApps
- List of RD Session Host servers per application
- Distribution of applications within RDS farm
- Any specific application needs, such as remote devices

User inventory RDS deployments are often measured in terms of the number of users supported. Taking count of users and how they access the RDS farm is a key requirement for RDS farm design. The user inventory should include the following:

- User count
- User connectivity (LAN, WAN, DSL, dial-up modem)
- User location (inside/outside firewall)
- Usage pattern

Business asset inventory A typical RDS farm has several components such as RDS server roles, applications, users, and machines, which can be deployed in variety of ways. Making sure the farm functions correctly and optimally requires you to keep track of several assets. This inventory item is critical not only to ensure compliance but also in the case of disaster recovery. This inventory list should include the following:

- Application license numbers
- RDS CALs
- RDS license server activation information
- Certificates for RDP file signing
- RDS Gateway authentication

Software inventory This inventory maintains details about all the software supporting each RDS server role as well as applications running in the RDS farm. This list should cover the following:

- OS version, patches, and service packs for both RDS server roles and clients
- RDS role–specific information such as OS configurations, installed applications (and versions)
- Content and location of RDS RAPs, RDS CAPs, and network access policies (if applicable) for RD Gateway
- Application prerequisites and dependencies

Hardware inventory The hardware inventory tracks details about the server hardware, the storage and network infrastructure, the peripheral devices, and the client devices used in the RDS farm. This list should include the following:

- Server models and capacity
- Network infrastructure
- Storage systems
- Peripheral devices such as remote printers
- Client hardware

Having this information handy will definitely ease your virtual RDS deployment. Equally important is the discipline to review and update this information on a regular basis to reflect any changes in your RDS deployment.

Checking Application Compatibility

Because RDS allows applications to be accessed by multiple users concurrently, you need to ensure that applications run correctly in such an environment. Most 32-bit and 64-bit applications should run without any changes on an RD Session Host server, but you need to pay special attention to any application compatibility issues. These issues typically fall into the following categories:

Application installation Single-user applications can have installation issues on a RD Session Host server if the install process creates files and registry entries that can be read by other user accounts logging on to the RD Session Host server.

Concurrent usage of applications Files, registry entries, pipes, IP addresses, and ports are used concurrently by multiple instances of applications when installed on an RD Session Host server. Applications should be written to ensure that any concurrent access/modification of these resources does not result in data corruption and other unexpected behavior.

User data privacy in a multiuser environment Applications running in a multiuser RD Session Host server environment must protect against any privacy breaches and unexpected sharing of data and user preferences.

Application performance Application performance when running multiple instances on an RD Session Host server should never be a bottleneck because of any network bandwidth or hardware constraints.

Remote devices in a Remote Desktop Services environment Application features and functionalities that rely on the clipboard, printers, and drives may not work over a remote session.

NOTE With the introduction of Windows Server 2008 R2, the Windows Server operating system is compatible only with 64-bit processor architectures. Although this helps the scalability of RD Session Host, it can create application compatibility issues. You need to be sure to address these issues as well.

Microsoft has launched an RDS application compatibility effort to ensure that applications run reliably on the RDS platform. It provides a runtime program analysis tool, RDS Application Analyzer (Figure 9.2), which can help you determine the compatibility of an application with an RD Session Host server before deploying it. The tool summarizes the incompatible behavior between the RD Session Host server and an application and provides recommendations for deploying the application on an RD Session Host server.

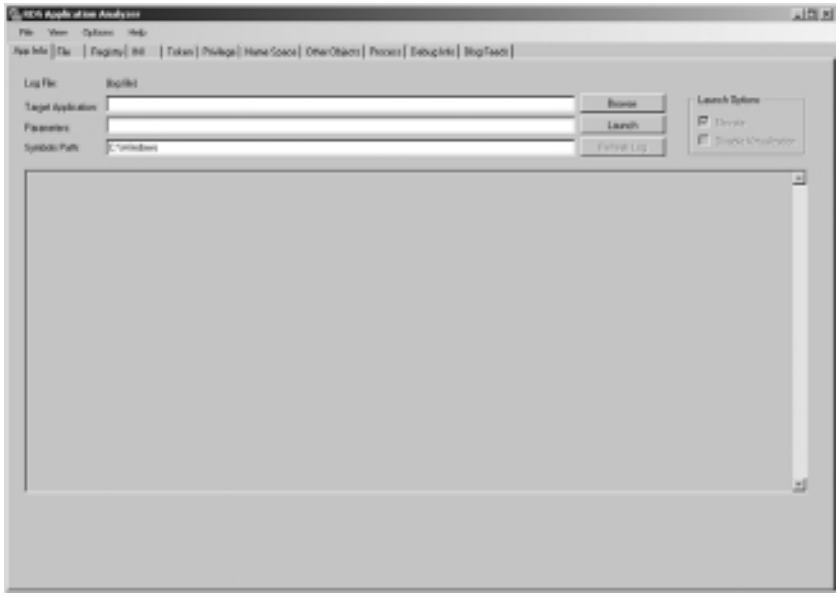

FIGURE 9.2
RDS Application Analyzer

The RDS Application Analyzer uses the Microsoft Application Verifier as the back end. When you analyze an application, the tool intercepts function calls from that application into the operating system and notes the calls and the parameters passed. Depending on the function call and the parameters passed, it decides whether this call is going to cause problems on an RD Session Host server and categorizes the results as warnings or problems.

It is important to understand that you will need to interpret the RDS Application Analyzer output to check whether the application is really compatible. The RDS Application Analyzer will flag a problem if it detects that the application has made an API call that is not compatible with RD Session Host. However, the call may just be part of checks that the application is trying to do.

Defining User Requirements

RDS deployments support a large number of users, with each of these users accessing applications in their own way. The end user experience often dictates the SLAs in place, so it is key to understand user requirements in an RDS farm. Several factors impact the user requirements and in turn the performance of your RDS deployment. These factors include the following:

- User count
- User location

- User software configuration
- User category
- User profile and data

In this section, we will discuss these factors so that you can design your RDS farm to provide a satisfactory end-user experience:

User count User count is the obvious factor that will determine the size and performance of your RDS environment. It is common to talk about RDS deployments in terms of total user count or number of users.

User location End users access the RDS environment from various types of locations. Not only can they be geographically dispersed, but they may be using LAN or WAN connections, may be inside a firewall or outside, and may be connected into the corporate network using VPN. Regardless of their location, the user expects a satisfactory RDS experience. Understanding user count at each location and their connectivity options will help you design the proper topology and size of the various RDS roles.

User software configuration Users accessing an RDS environment will likely have different client software configurations. For example, the operating system they are using will determine the versions of RDC and RDP used for connecting to the RDS farm. This will affect the amount of data that is communicated back and forth from the client to the RDS farm. Other factors such as web browsers and their add-ons, authentication requirements, certificates, and other signed files also will influence client software configuration.

User category Users in an RDS environment use their applications in a variety of ways. The actual usage of an application—the sequence of actions, the frequency and duration, the application features—can vary from one user to other. To make things simple, it is common practice to place users into one of three categories, based on their general usage behavior in their applications: Heavy, Normal, or Light. You may also read about other types of user categories such as knowledge workers, task workers, and data entry workers. The key idea is to classify users in fewer, reasonably well-defined groups that can help in your capacity planning and sizing exercises.

User profiles and data Multiple users and RD Session Host servers in an RDS farm typically call for a well-defined strategy for user profiles and user-specific data. Whether you use roaming profiles or mandatory profiles, you need to determine how much space is required for user profiles and where to store the user profiles. Although user profiles are generally quite small, the storage requirements do add up as the user count goes up. If users are allowed to store data in a My Documents folder on the storage server, then the storage space required can be significant. You can make a rough estimate for the total user data storage space required by multiplying the number of users and the maximum storage space each user needs. Several best practices such as the following need to be put into practice based on these requirements:

- Separating the user data (in My Documents and Application Data files) from the profile
- Storing user profiles and data on separate network drive or file server
- Limiting a user profile size using quotas
- Ensuring that roaming user profiles are cached locally on the RD Session Host server

Sizing Your RDS Deployment

Sizing an RDS deployment can be a challenging exercise. Several factors will determine the workload that an RDS farm can support. First, there is a lot of flexibility when deploying in the farm mode, and you can combine various RD server roles to change the workload on a given server. Second, the workload of these RDS servers depends upon the wide variety of applications that may be running on them and the different ways in which users work with these applications. Unlike other tier-1 applications such as database, email, and web servers that usually exhibit a limited type of workload transactions, the workload in an RDS environment can vary drastically. Furthermore, even the same set of applications can be used in so many ways that estimating the user count can be a tricky exercise. We have come across RDS deployments where the user count ranges anywhere between 10 users to 150 users per RD Session Host server. Fortunately, all the ancillary roles in an RDS farm, except the RD Session Host role, are usually lightweight roles. So, you must make sure that none of the ancillary roles creates a bottleneck in your RDS farm. As suggested in previous sections, we recommend you start with a single CPU and a 2 GB virtual machine for these ancillary RDS roles and scale out using multiple such virtual machines if needed. Not only will such horizontal scaling support more users, but it will also provide fault tolerance to your RDS farm. You will need to adjust your user count to accommodate any capacity headroom requirements you may have.

The RD Session Host is the one that runs multiple application instances associated with the numerous users. Because the performance is subject to so many factors, there is no easy way to size RD Session Host servers. To be able to estimate the needed size of your RDS farm with reasonable accuracy, you will need to test your deployment using a realistic workload that represents your environment. Often we have seen customers set their expectations based on published high watermark benchmarks. Unfortunately, such benchmark reports typically only use Microsoft Office applications and do not take into account all the user requirements we discussed in previous sections. You can approach the sizing exercise in two ways:

Understanding resource requirements Understand the resource requirements (CPU, RAM, disk, network) for individual applications in your environment for a single user. You can then extrapolate this number based on the user count you plan to support.

Deploying a proof of concept Deploy a proof of concept based on the previous extrapolations, and simulate user workload in this pilot environment. You can either generate a synthetic load based on some load-generating products or decide to have real-life users access this proof-of-concept environment.

The reliability of this sizing exercise will depend upon how accurately you are simulating the actual usage scenario. Any change in hardware and software configuration, applications used, specific usage of each application, or user-related details will affect the deployment scenario and make the results less accurate. You will need to be particularly disciplined to follow performance and capacity planning best practices such as establishing correct baselines and avoiding an "apples to oranges" comparison.

TIP The Microsoft RDS team recently published a white paper called "Remote Desktop Session Host Capacity Planning in Windows Server 2008 R2" that covers this topic. You can access this white paper from the RDS team blog at `http://blogs.msdn.com/rds/archive/2010/02/26/remote-desktop-session-host-capacity-planning-in-windows-server-2008-r2.aspx`.

During these capacity planning tests, you can adjust the user load in the RDS farm as needed while monitoring the user feedback. Make sure you use vSphere tools such as `esxtop` and vCenter to monitor performance for the virtual environment. You can use the standard Windows Performance Monitor to monitor the RDS-specific performance counters to determine any potential bottlenecks. You can then adjust the software and hardware configurations in your proof of concept to overcome any limiting factors you uncover during the testing. The data obtained from these tests will help you size the memory, processor, disk I/O, and network resources that will be required to deliver a satisfactory end-user experience in your RDS farm. You will need to adjust these results to take into account any capacity headroom you want to put in place to accommodate any usage spikes.

Deployment Considerations

RDS farms comprise different server roles, often deployed in a horizontal scale-out topology. Given the number of server instances, applications, and multiple configuration settings, customers want to know whether there are guidelines to follow when deploying in a vSphere environment. Customer questions range from basic virtual machine creation to managing these virtual machines on an everyday basis.

vSphere definitely provides you with opportunities to simplify your virtualized RDS farm deployments, compared to deployments in physical environment. VMware features such as vSphere DRS, HA, and templates and clones make it easy to address these deployment questions. When deploying and managing an RDS farm, the administrative and management flexibilities that virtualization offers often help you get over any objections you may have for RDS virtualization. In this section, we will discuss the key deployment considerations for a virtual RDS farm:

- Initial creation of RDS virtual machines
- Ongoing provisioning strategy
- vSphere DRS and HA considerations

Initial Creation of RDS Virtual Machines

When you create the various RDS virtual machines, you can either use the VMware Converter tool or create fresh virtual machines leveraging vSphere's templates and clones.

Physical to virtual (P2V) migration using VMware Converter is usually more convenient and makes it faster to create VMs. However, P2V can result in two key issues for RD Session Host server: larger VMs and extra processes within a VM. Both can have a serious impact on VM performance.

Most server-class machines today have two sockets or more with at least a single-core processor in each socket. The newer machines have multicore processors. Add to this the fact that almost all physical servers are over-provisioned, with average CPU utilization at best in the teens. Consequently, P2V invariably ends up creating larger VMs. Larger VMs means SMP-related virtualization overheads and reduced DRS effectiveness. Also, physical servers typically have administration tasks and services such as SNMP monitoring and backup agents running in the background. These additional services can result in several performance issues such as IRQ conflicts and polling/interrupt-driven processes, all of which will typically exaggerate virtualization overhead after P2V.

RDS environments are especially sensitive to larger VMs and extra processes within a VM. Most RDS roles are lightweight services and do not need larger VMs. RD Session Host servers already exhibit a high number of context switches between various processes, lots of process creation and destruction, and high paging activity. These two issues in P2V'd VMs—SMP-related virtualization overheads and having additional processes running in a P2V'd VM—just aggravates this further. Hence, the best-practice recommendation is to start with fresh VMs.

In real life, you may not be able to avoid using VMware Converter in some cases. For example, you may have an old Windows server with some legacy OS/application stack for which you cannot do a fresh install, recompile, or port to a newer OS. Thus, P2V is the best available option. In that case, consider the following things:

- Before you power on the new virtual machine after P2V, edit the VM settings to reduce the CPU count and memory, adjust the Windows HAL to the UP version if needed, and remove any other peripherals not required.
- Before P2V or after the P2V'd VM is powered on, turn off and disable all unneeded processes and services.

Ongoing Provisioning Strategy

RDS farms comprise several servers with different RDS roles and RD Session Host servers with a variety of applications. Scaling out with multiple virtual machines is recommended for scalability, load balancing, and high availability.

Maintaining a large number of servers in an RDS farm can be time-consuming. You can take advantage of templates and cloning features in VMware vSphere when provisioning RDS virtual machines. We discussed this in detail in Chapter 3. Using templates, you can build base templates that have the required Windows Server OS installed. You can then customize these templates further, based on RDS roles and applications, and lock down their configurations as needed.

You can easily clone or deploy these templates over and over again to reduce the time it takes to provision the various RDS roles. You can create a template with all the RDS ancillary roles installed and use it to deploy only the specific RDS service you needed. For RD Licensing and RD Session Host roles that save licensing and application information locally, respectively, you will need to create separate templates.

For RDS farms with RD Session Host servers that run different applications, you will need to maintain separate templates for each application group. If you elect to allow users access to the entire remote desktop, you may need only one template for each RD Session Host silo. However, if you elect to selectively publish applications using RemoteApp, you should consider creating a separate template for each RemoteApp in your provisioning process. Otherwise, you will need to configure the RemoteApp separately after deploying the RD Session Host template.

When you have multiple images (Gold, Silver, and so on), you will need a plan to deploy your templates efficiently. Make sure to keep the templates up-to-date with Microsoft updates and patches. You can leverage VMware Update Manager to update your templates and VMs. If you have an application update or change, update that silo's template and deploy the template. This will keep all your RD Session Host servers in an RDS farm identical and save any regression testing cycles.

vSphere DRS and HA Considerations

Several server roles in RDS farm use a horizontal scale-out strategy for scalability and fault tolerance. For RDS deployments in vSphere environment, we recommend you utilize vSphere

DRS to load-balance your RDS servers in a vSphere cluster. You can set up a cluster specifically for your RD Session Host servers if needed. Set DRS Automation Level to Fully Automated, and set Migration Threshold to Apply Recommendations With Four Or More Stars. You should also configure the DRS anti-affinity rules so that redundant servers are not on the same vSphere host at the same time to prevent single points of failure.

We recommend you leverage vSphere HA to address any unplanned vSphere host downtimes. With VMware HA, a virtual machine in an RDS farm will automatically restart on another vSphere host. Depending upon your SLAs, this can provide you with enough high availability without the complexity associated with Windows Failover Clustering. You can extend this high availability further using VMware FT, which will even avoid the downtime associated with restarting a virtual machine.

Virtualizing Remote Desktop Services Proof of Concept

Source Discovery employees are dispersed geographically and frequently need to make corporate information easily available to workers from remote locations and need to ensure that important internal company information is as secure as possible. The employees often need to access internal company resources from home or from a remote location. These resources mostly include Microsoft Office documents, presentations, and spreadsheets. Source Discovery plans to deploy an RDS environment to address this need. Given the recent success of virtualizing several Microsoft tier-1 server applications using VMware vSphere, Source Discovery wants to virtualize all RDS roles for this proof of concept.

Source Discovery anticipates deploying an RDS farm with at least two RD Session Host servers to meet the user capacity and anticipates using RD Connection Broker for user session load balancing. Because users may or may not be logged to the corporate network using a virtual private network (VPN) connection, Source Discovery plans to deploy both the RD Gateway and RD Web Access roles.

The company also wants to make this proof-of-concept environment highly available for all RDS roles wherever possible. During the proof of concept, the company will use Windows Network Load Balancing to provide failover and redundancy for the RD Gateway. When Source Discovery transitions the farm to production, it will use a hardware load balancer to provide load distribution, redundancy, and availability for the external users. Source Discovery will also extend the fault tolerance of this environment using vSphere DRS, HA, and FT.

You will use Windows Server 2008 R2 templates from Chapter 3 to build the various RDS role virtual machines. As we show how to deploy these servers, you will configure the vCPU, memory, disks, and networking appropriately to match the application requirements. For example, you will start with one vCPU and 2 GB of memory for the ancillary roles in the RDS farm, except RD Session Host. If needed, you will increase this to two vCPUs and 4 GB of memory.

We're assuming you have a thorough knowledge of the RDS technology product, so we will not cover every configuration step in this proof of concept. In the next sections, we will cover the configuration for each server in the RDS farm. We will also discuss the high-availability setup if applicable.

TIP Microsoft has published several deployment guides for the different RDS roles in an RDS server farm. You can leverage these guides for deployment in the vSphere environment as well. You can access this information from the RDS TechNet website at `http://technet.microsoft.com/en-us/windowsserver/ee236407.aspx`.

Configuring the Active Directory Virtual Machine

The RDS farm needs an Active Directory server to validate whether a user has the right to connect to the RD Session Host server, to manage group policies that configure RD Session Host and the user sessions running on them, and to maintain per-user RDS CAL consumption. When using per-user licensing, you have to make sure that the RD Licensing server's domain is trusted by the user account domains; otherwise, the RD Licensing server will not be able to issue licenses. You can use Active Directory with other RDS roles in one machine, but best practices recommend you maintain separate virtual machines for each role to avoid any service interruptions.

In this proof of concept, you will leverage the Active Directory virtual machine that you deployed in Chapter 4. If you have already transitioned this virtual machine to production, you will need to use a clone of this virtual Active Directory deployment in a fenced-off network to avoid any conflicts.

Configuring the RD Licensing Virtual Machine

You will start the RDS proof of concept by deploying the RD Licensing virtual machine. You deploy this virtual machine from the Windows Server 2008 R2 template you created in Chapter 3. The virtual machine consists of the following hardware:

- One virtual CPU (vCPU)
- 2 GB of virtual machine memory
- One virtual network interface card (We have configured this card to use VMXNET3 driver.)
- One virtual disk for the operating system and RD Licensing role

Figure 9.3 shows the configuration of the RD Licensing virtual machine.

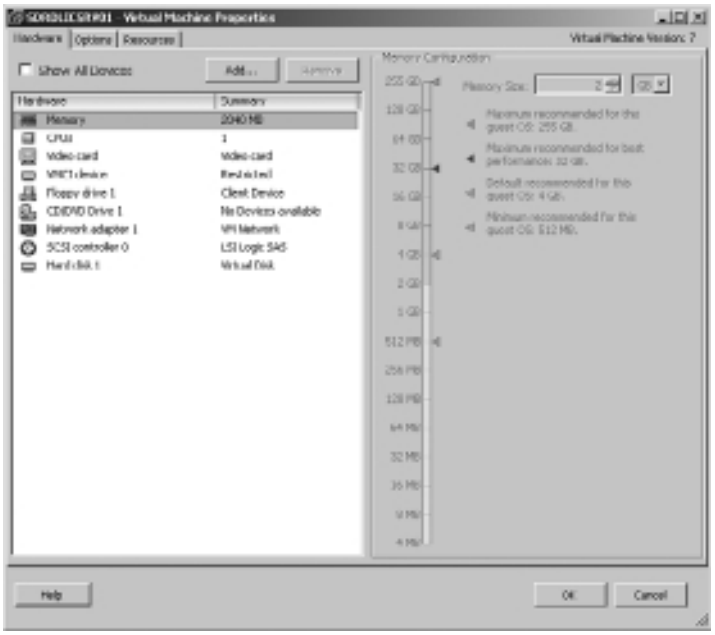

FIGURE 9.3
RD Licensing virtual machine

Next you install all the required operating service packs and critical updates and join the virtual machine to the Source Discovery domain before adding the RD Licensing server role. You can use the Add Roles Wizard in Server Manager to install the RD Licensing role service. Follow the steps in the wizard. On the Select Server Roles page, select the Remote Desktop Services check box, and then click Next (Figure 9.4).

Select the Remote Desktop Licensing check box on the Select Role Services page, and then click Next (Figure 9.5).

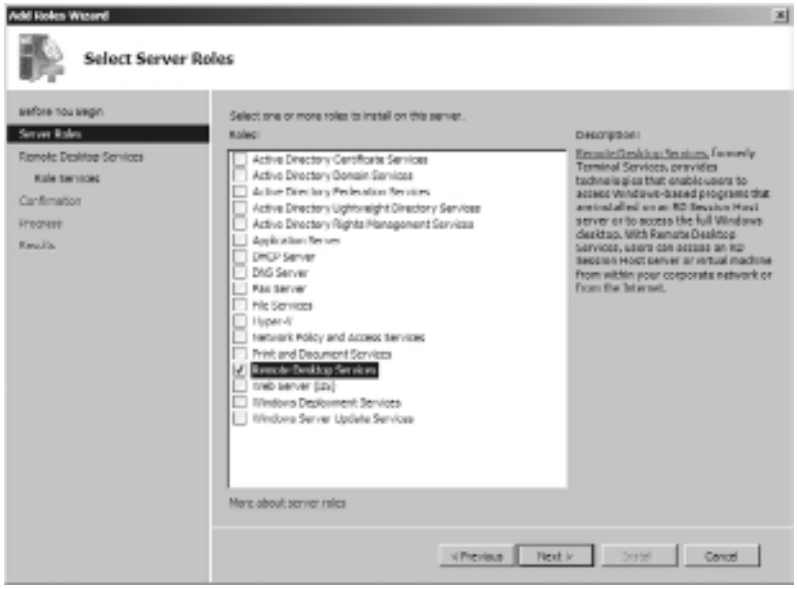

FIGURE 9.4
Remote Desktop Services role

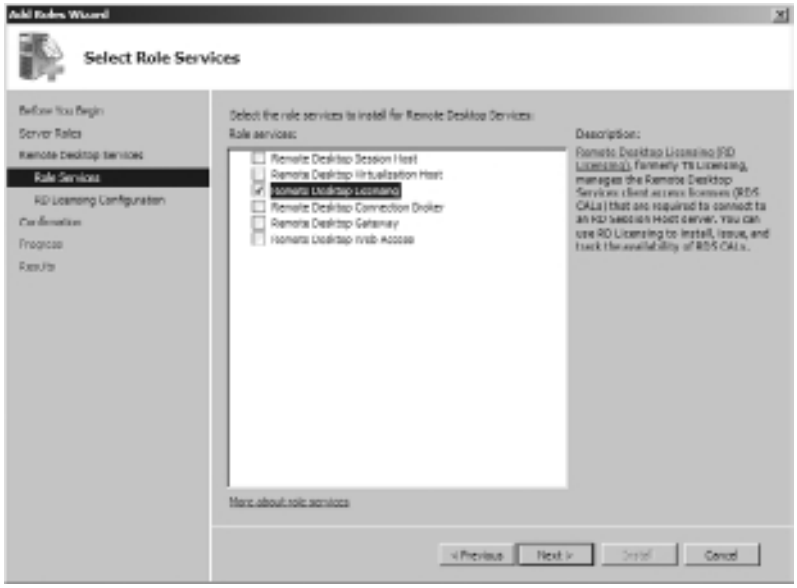

FIGURE 9.5
Remote Desktop Licensing role service

On the Configure Discovery Scope For RD Licensing page, select the Configure A Discovery Scope For This License Server check box, click The Forest, and then click Next (Figure 9.6). When installing the license server in a domain, you can use the forest scope to make the server most easily discoverable. You can restrict access to the license server to only those RD Session Host servers that need it.

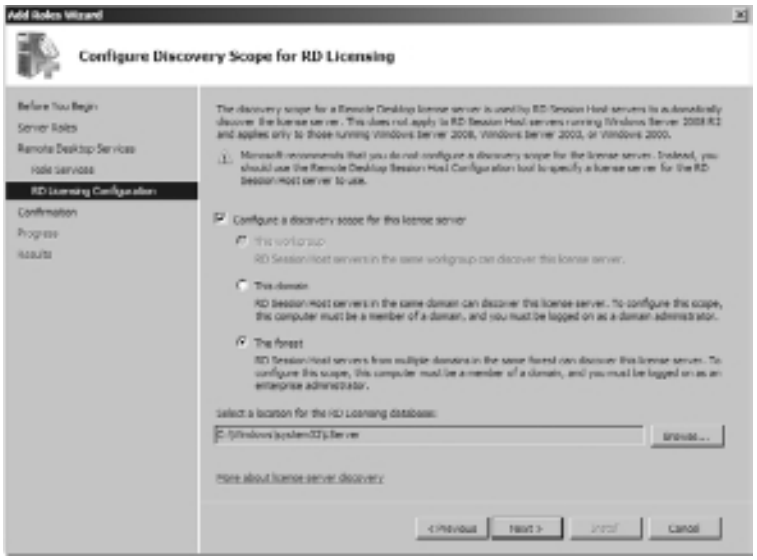

FIGURE 9.6
Configure Discovery Scope For RD Licensing

On the Confirm Installation Selections page, verify that the RD Licensing role service will be installed, and then click Install (Figure 9.7). This will start the RD Licensing role installation process. The Installation Progress page will show the installation progress.

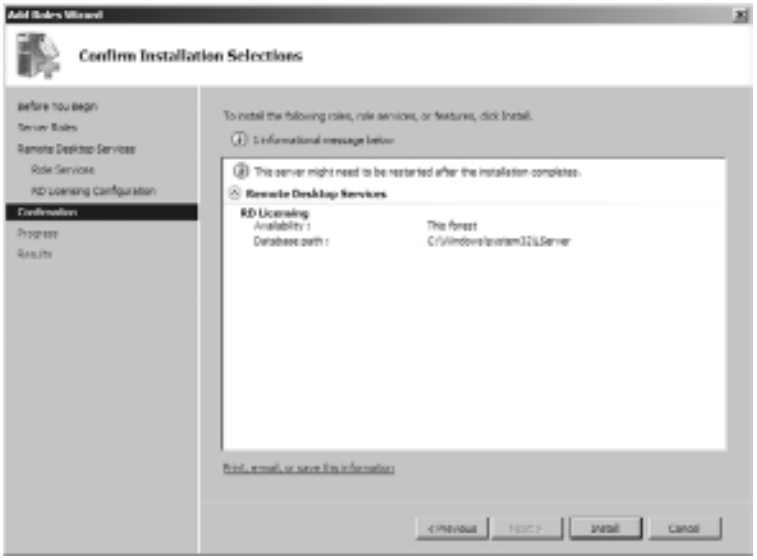

FIGURE 9.7
Confirm installation selections

On the Installation Results page, confirm that installation of the RD Licensing role service was successful, and then click Close (Figure 9.8).

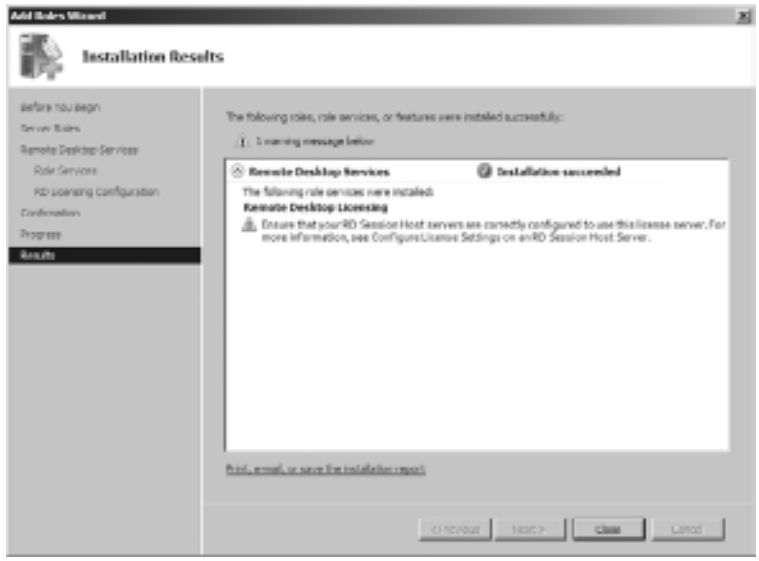

FIGURE 9.8
Installation Results

For activating the RD Licensing server and issuing RDS CALs, the Microsoft Clearinghouse uses OS-specific information on the RD license server. This makes it difficult to move the RD Licensing server role to a different machine in a physical environment. In a virtual environment, this is no longer an issue because you can leverage vSphere vMotion and HA to make the RD Licensing server highly available. In this proof of concept, we will use the 120-day grace period that Microsoft provides, but we will document the steps to activate the RD Licensing server and to install RDS CALs.

You need Remote Desktop Licensing Manager to activate an RD Licensing server. By default it needs Internet access to complete the activation automatically. You can start the Activate Server Wizard by right-clicking the license server name (Figure 9.9). Follow the steps in the wizard to provide the necessary company information.

When the RD license server is activated, you need to install RDS CALs. From Remote Desktop Licensing Manager, right-click the license server to start the Install Licenses Wizard. You need to select the appropriate program through which you purchased your RDS CALs. Based on the license program that you select, the wizard determines what information you need to provide. In most cases, you must provide either a license code or an agreement number. After you enter the required information, you must select the appropriate product version, license type, and quantity of RDS CALs for your environment based on your RDS CAL purchase agreement. The wizard contacts the Microsoft Clearinghouse automatically and processes your request. The RDS CALs are then automatically installed onto the license server. At the end of the process, RD Licensing server is ready to issue RDS CALs to clients that connect to an RD Session Host server.

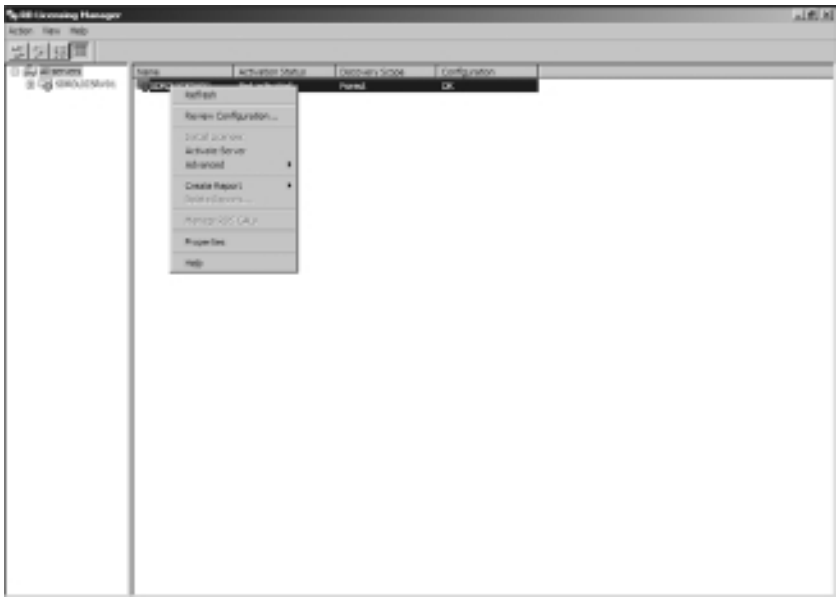

FIGURE 9.9
RD Licensing Manager

You can also review the configuration from the RD Licensing Manager window. Use the licensing diagnostics tool to validate any RD Licensing–specific information and details about any problems you face (Figure 9.10).

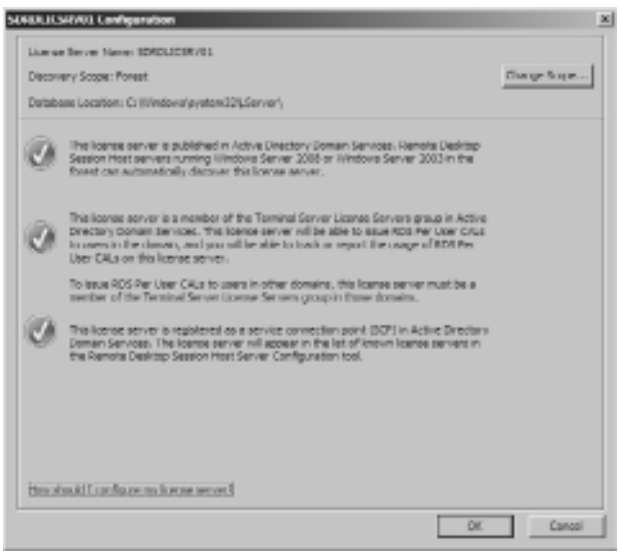

FIGURE 9.10
RD Licensing server configuration review

The RD Session Host server needs to contact the RD Licensing server to access RDS CALs. For best availability, it is recommended you have multiple RD Licensing servers and split the RDS CAL license packs among them. This way you avoid the possibility of service outage because an RD Session Host server cannot reach the single RD Licensing server in your RDS farm. For this proof of concept, you will deploy a second RD Licensing server following the previous steps. Figure 9.11 shows the screenshot with two RD Licensing servers in the proof-of-concept RDS server farm. Later you will also configure RD Session Host servers that point to these two RD Licensing servers.

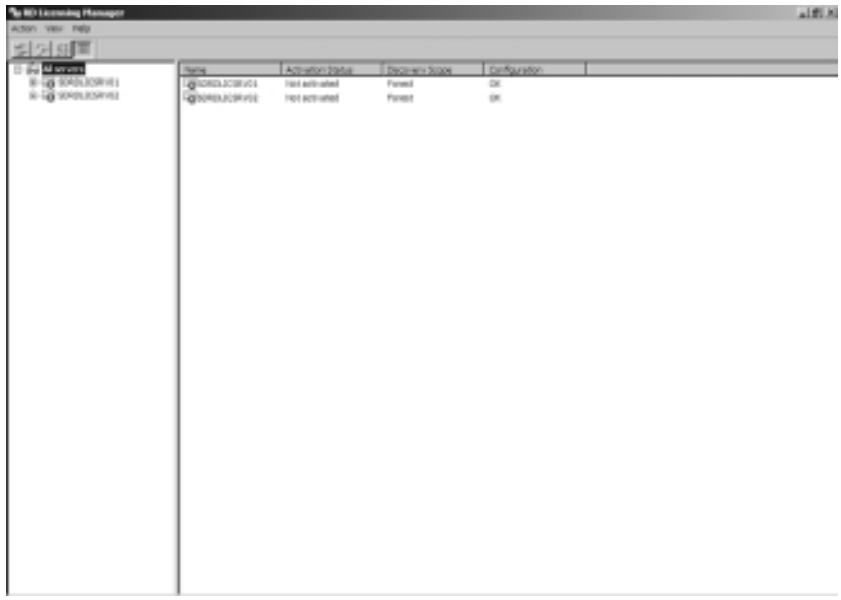

FIGURE 9.11 Two RD Licensing servers in the proof-of-concept deployment

Configuring the RD Session Host Role

RD Session Host servers typically are resource intensive because they support multiple users, who may be doing various kinds of actions with different types of applications. Therefore, you should dedicate a virtual machine only for the RD Session Host role and not combine with any other roles in an RDS farm. Because the RD Session Host server performance is subject to several factors, it is difficult to predict what hardware you should configure for the RD Session Host virtual machine to support your requirements. In general, you should give this role as many resources as you can; otherwise, it will result in the performance degradation of the RDS farm. Here are some specific guidelines:

CPU We recommend starting with two vCPUs for the RD Session virtual machine. Our experience shows that having the additional vCPUs results in much better end-user response times. Several other factors such as SMP scalability factor, user count spikes, and application CPU usage variability will need to be considered as well. In production environments, it is common to leave some capacity headroom for unexpected usage spikes. You can configure a four-vCPU virtual machine, but we recommend using a building-block approach by staying at a two-vCPU virtual machine and instead adding more RD Session Host virtual machines to your RDS farm.

Memory In general, load up the RD Session Host virtual machine on memory as much as you can. Start with 4 GB for the 32-bit Windows Server 2008 version and 8 GB (or 16 GB recommended) for the 64-bit versions. Setting memory reservation for the RD Session Host virtual machine will prevent vSphere from doing any swapping.

Storage RD Session Hosts stress the storage subsystem for the following I/O:

- The OS and application files
- The page file
- User-specific data such as user profile, My Documents, application data, temporary file folders, and so on

You should plan for enough physical spindles, utilize the fastest disks possible, and use RAID 1 or 10 for performance and redundancy. Utilize separate thick-provisioned virtual disks carved out on separate physical volumes for your operating system, applications, user-specific data, and page file. Using folder redirection is recommended to improve user data I/O performance. Follow the best recommendations to size the page file three times the size of the virtual machine memory.

Network We recommend using a multihome server with one network for RDP traffic and the other for file and print requests. Network bandwidth depends upon different settings for the RDP session, such as pixel dimensions, color depth, redirected devices, and compression schemes. Graphic-rich applications will need higher network bandwidth. You must also account for network traffic if user profiles are not stored locally.

TIP Microsoft has published performance-tuning guidelines for RD Session Host in the following white paper: `www.microsoft.com/whdc/system/sysperf/Perf_tun_srv-R2.mspx`.

You will leverage the Windows Server 2008 R2 template as discussed in Chapter 3 to build the RD Session Host virtual machine for this proof of concept. Start with the following resources for the virtual machine:

- Two virtual CPUs (vCPUs).
- 8 GB of virtual machine memory.
- One virtual network interface card (We have configured this card to use VMXNET3 driver.)
- Two virtual disks: one virtual disk is provisioned for the Windows operating system, Microsoft Office applications, and user profiles. The second virtual disk is provisioned for the page file.

Figure 9.12 shows the configuration of the RD Session Host virtual machine.

You will create an RDS farm with two RD Session Host servers to provide load balancing and redundancy. So, you duplicate this virtual machine configuration for the second RD Session Host server. You will leverage vSphere HA to further extend the availability for each RD Session Host virtual machine.

FIGURE 9.12
RD Session Host server VM configuration

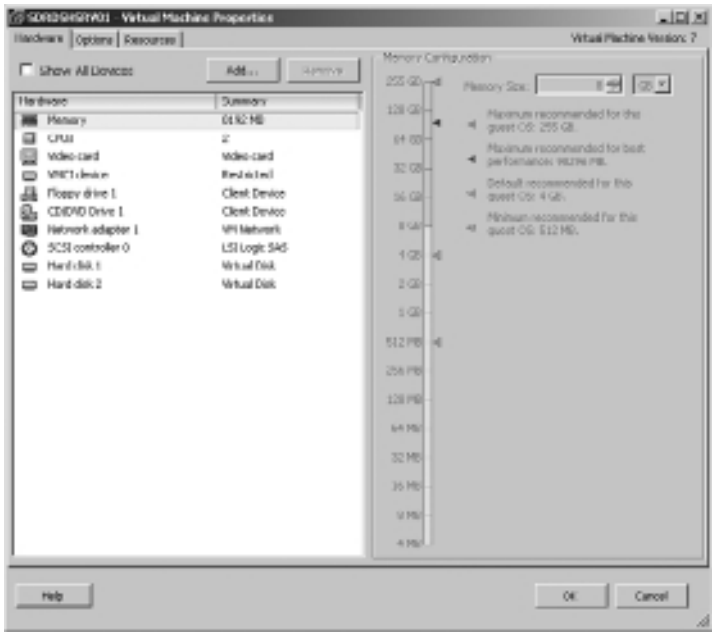

The RD Session Host should be as generic as possible; specifically, all load-balanced servers should contain the same applications so the end user sees no difference. Also, they should never contain any user-specific data at all. In this proof of concept for Source Discovery, initially you will run only Microsoft Office applications on both RD Session Hosts.

Next you install all the required operating service packs and critical updates and join the virtual machine to the Source Discovery domain. You move the paging file to the second virtual disk and size it three times the size of the memory configured for the RD Session Host virtual machine (Figure 9.13).

FIGURE 9.13
RD Session Host paging file location

Next you use the Add Roles Wizard in Server Manager to install the RD Session Host role service. Follow the steps in the wizard. Make sure you select the Remote Desktop Session Host check box on the Select Role Services page (Figure 9.14).

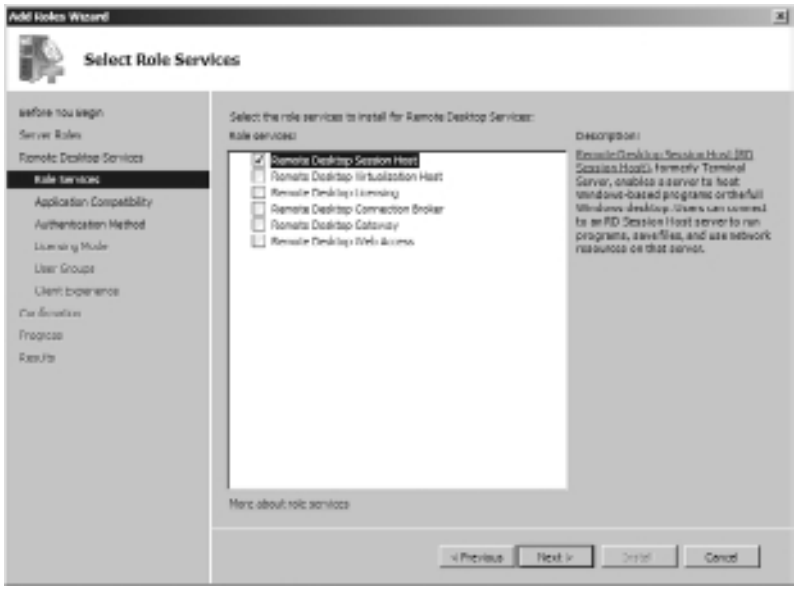
Add Roles Wizard
Select Role Services
Before You Begin
Server Roles
Remote Desktop Services
Role Services
Application Compatibility
Authentication Method
Licensing Mode
User Groups
Client Experience
Confirmation
Progress
Results
Select the role services to install for Remote Desktop Services:
Role services:
Remote Desktop Session Host
Remote Desktop Virtualization Host
Remote Desktop Licensing
Remote Desktop Connection Broker
Remote Desktop Gateway
Remote Desktop Web Access
Description:
Remote Desktop Session Host (RD Session Host), formerly Terminal Server, enables a server to host Windows-based programs or the full Windows desktop. Users can connect to an RD Session Host server to run programs, save files, and use network resources on that server.
More about role services
< Previous
Next >
Install
Cancel

FIGURE 9.14
RD Session Host role service

Click Next on the Uninstall And Reinstall Applications For Compatibility page (Figure 9.15).

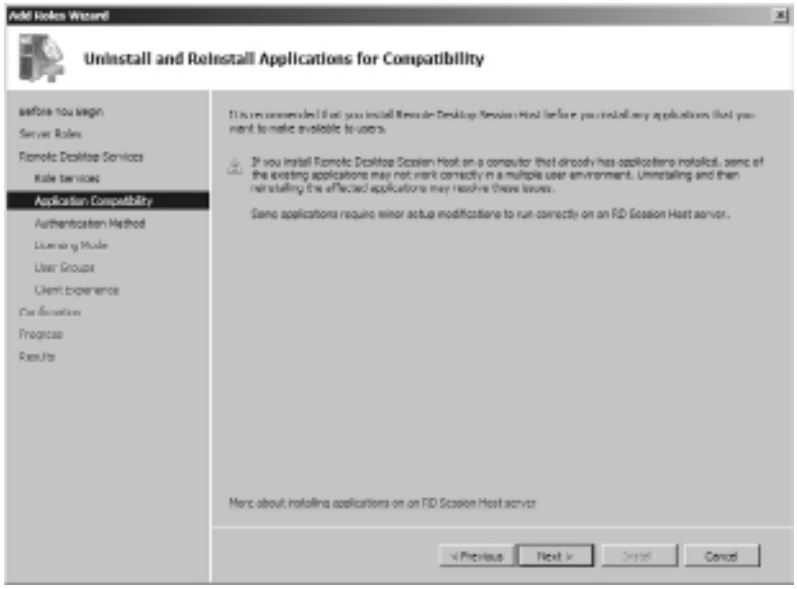
Add Roles Wizard
Uninstall and Reinstall Applications for Compatibility
Before You Begin
Server Roles
Remote Desktop Services
Role Services
Application Compatibility
Authentication Method
Licensing Mode
User Groups
Client Experience
Confirmation
Progress
Results
If you install Remote Desktop Session Host on a computer that already has applications installed, some of the existing applications may not work correctly in a multiple user environment. Uninstalling and then reinstalling the affected applications may resolve these issues.
Some applications require minor setup modifications to run correctly on an RD Session Host server.
More about installing applications on an RD Session Host server
< Previous
Next >
Install
Cancel

FIGURE 9.15
Application compatibility warning

On the Specify Authentication Method For Remote Desktop Session Host page (Figure 9.16), choose Do Not Require Network Level Authentication since Source Discovery needs to support Windows XP SP2 client computers. If you require Network Level Authentication, only clients supporting CredSSP (Windows Vista or XP SP3) will be able to connect.

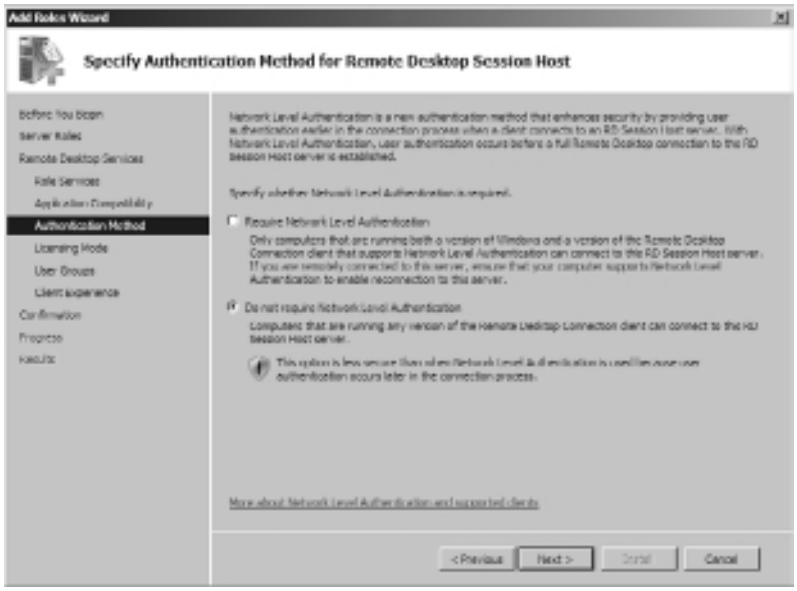

FIGURE 9.16
Authentication method for RD Session Host

For the proof of concept, choose to configure the licensing mode later on the Specify Licensing Mode page (Figure 9.17).

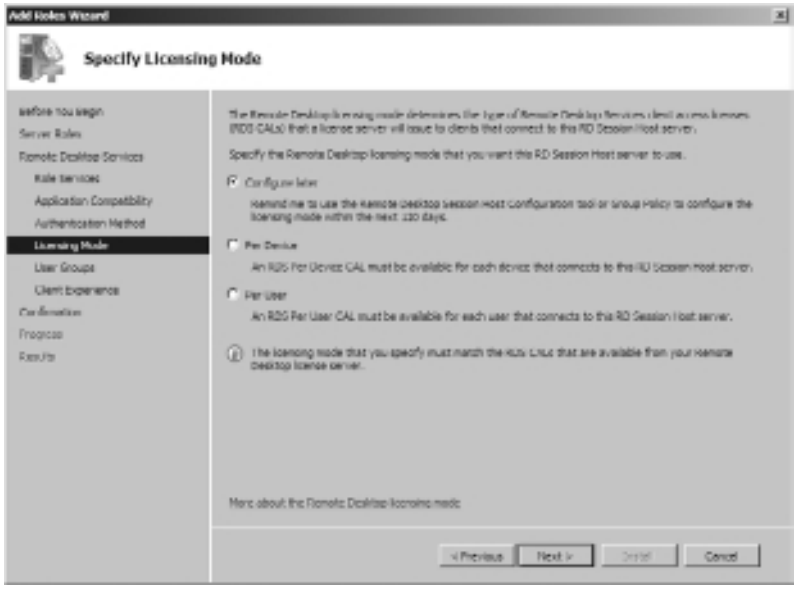
Add Roles Wizard
Specify Licensing Mode
Before You Begin
Server Roles
Remote Desktop Services
Role Services
Application Compatibility
Authentication Method
Licensing Mode
User Groups
Client Experience
Confirmation
Progress
Results
Specify the Remote Desktop licensing mode that you want this RD Session Host server to use.
Configure later
Remind me to use the Remote Desktop Session Host configuration tool or Group Policy to configure the licensing mode within the next 120 days.
Per Device
Per User
An RDS Per User CAL must be available for each user that connects to this RD Session Host server.
< Previous
Next >
Install
Cancel

FIGURE 9.17
RD Session Host licensing mode

Next you select User Groups Allowed Access To This Remote Desktop Session Host Server. Click Next on the Configure Client Experience page. On the Confirm Installation Selections page, verify that the RD Session Host role service will be installed, and then click Install (Figure 9.18). This will start the RD Session role installation process.

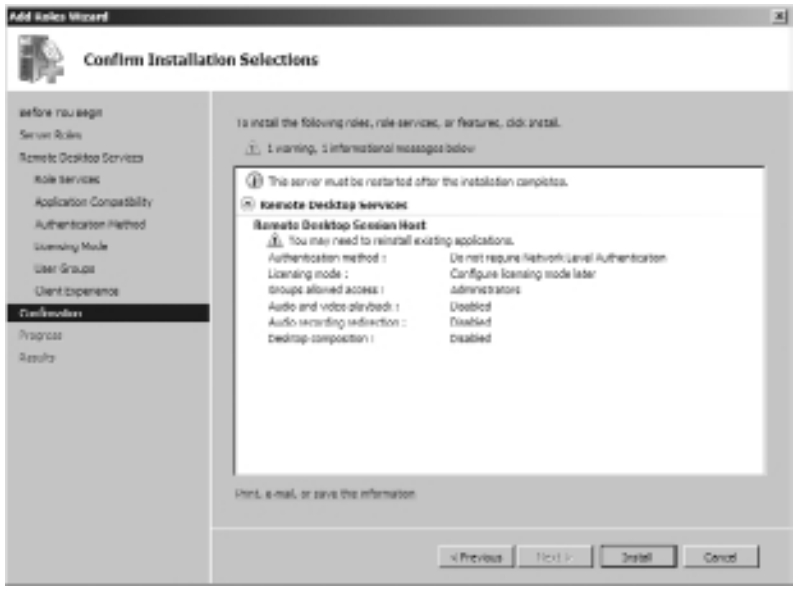

FIGURE 9.18
Confirming installation selections

On the Installation Results page, you are prompted to restart the server to finish the installation process (Figure 9.19). Click Close, and then click Yes to restart the server.

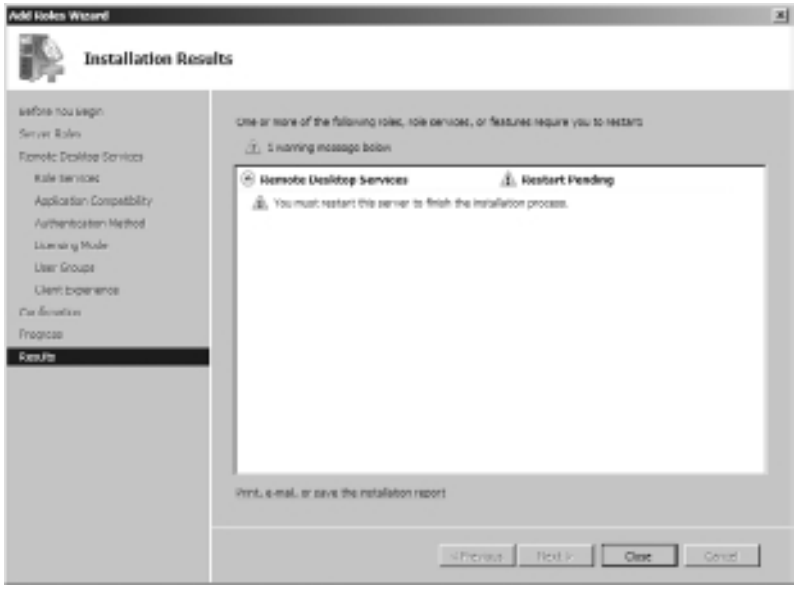

FIGURE 9.19
Installation Results

The installation process will continue after the server restarts (Figure 9.20).

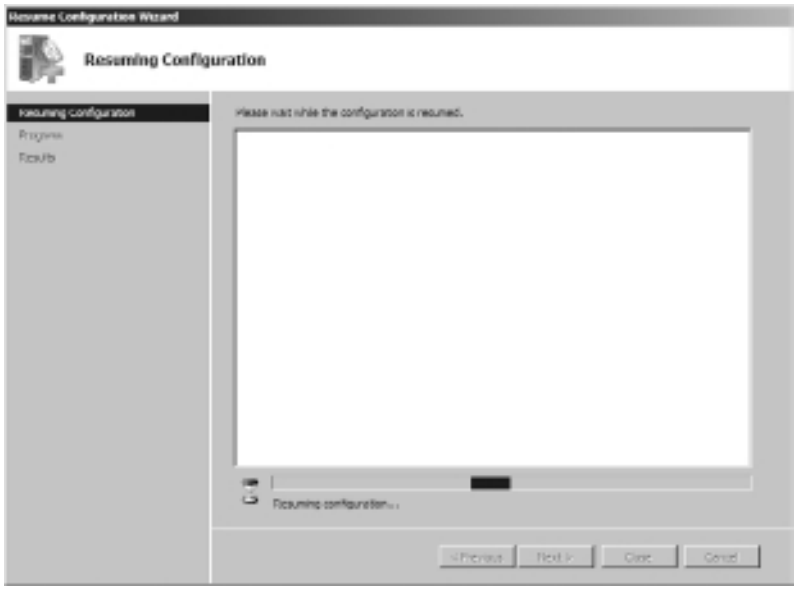

FIGURE 9.20
Resuming configuration

When the installation finishes, the Installation Results page should show that the RD Session Host role service installation has succeeded (Figure 9.21).

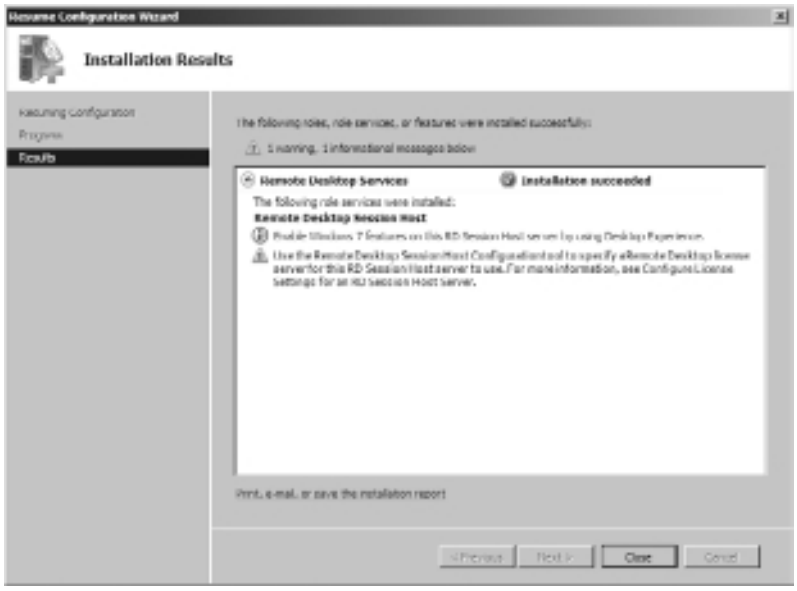

FIGURE 9.21
RD Session Host installation results

Next you need to configure RD Session Host server to ensure that it can contact the RD Licensing server to request RDS CALs for clients. The RD Session Host Configuration tool shows the current configuration of the RD Session Host server (Figure 9.22). You can use this tool to configure the licensing mode, specify multiple RD Licensing servers, and configure the RD Session Host server farm.

FIGURE 9.22
RD Session Host Configuration tool

To configure the Remote Desktop licensing mode, double-click Remote Desktop Licensing Mode under Licensing. You can specify either Per Device or Per User, depending on which is most appropriate for your environment (Figure 9.23).

FIGURE 9.23
RD Session Host configuration – licensing mode

From the Licensing tab, you can also add specific license servers for the RD Session Host server to use. The Add License Server dialog box should be already populated with known license servers. For the proof of concept, select both the license servers, and add them to the list of specified license servers (Figure 9.24).

FIGURE 9.24
RD Session Host configuration – specified license servers

Figure 9.25 shows the licensing information configured at the end of these steps.

FIGURE 9.25
RD Session Host configuration – licensing configuration

Make sure you add necessary users to the Remote Desktop Users group so they log on to the RD Session Host server remotely (Figure 9.26).

You should run the Best Practices Analyzer from Server Manager to scan the specific RDS role and ensure that you are complying with the best practices in your deployment (Figure 9.27). The RDS Best Practices Analyzer, introduced in Windows Server 2008 R2, scans the different RDS server roles to report any violations in RDS best practices. Currently it checks for certain configuration settings to make sure that RDS Server roles are configured correctly and also performs a few operational checks to make sure those roles can communicate with each other as required.

FIGURE 9.26
Remote Desktop Users group on RD Session Host

FIGURE 9.27
RDS Best Practices Analyzer

You are now ready to install applications on the RD Session Host. For this proof of concept, you install Microsoft Office 2007 on each RD Session server. Once the RD Session Host role is configured, you will notice a new shortcut under Control Panel (Figure 9.28).

Programs
Control Panel • Programs •
Search Control Panel
Control Panel Home
System and Security
Network and Internet
Hardware
Programs
User Accounts
Appearance
Clock, Language, and Region
Ease of Access
Programs and Features
Uninstall a program | Turn Windows features on or off | View installed updates
Run programs made for previous versions of Windows | How to install a program
Default Programs
Change default settings for media or devices | Make a file type always open in a specific program
Set your default programs
Install Application on Remote Desktop...
VMware Tools

FIGURE 9.28
Installing applications on RD Session Host

Follow similar steps to deploy and configure a second RD Session Host server for the proof-of-concept RDS farm. In large RDS farm environments, you should configure the RD Session Host servers using Group Policy, not the RD Session Host configuration tool, to avoid any inconsistencies among RDS servers. Figure 9.29 shows RD Session Host group policy specifying the RD Licensing servers in the proof-of-concept deployment. In virtual environments, using proper templates and cloning from them can alleviate this issue during the initial installation.

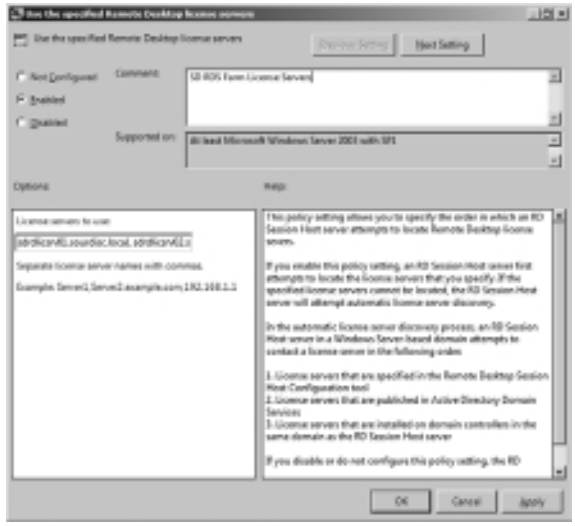
Use the specified Remote Desktop license servers
Previous Setting
Next Setting
Not Configured
Enabled
Disabled
Comment:
Supported on:
Options:
Help:
License servers to use:
Separate license server names with comma.
This policy setting allows you to specify the order in which an RD Session Host server attempts to locate Remote Desktop license servers.
If you disable or do not configure this policy setting, the RD
OK
Cancel
Apply

FIGURE 9.29
RD Session Host group policy for RD Licensing servers

Configuring the RD Web Access Virtual Machine

To enable end users to access remote desktops and RemoteApps from a web browser, you need to deploy the RD Web Access virtual machine. As with other RDS roles, you start with the Windows Server 2008 R2 template you created in Chapter 3. The RD Web Access virtual machine is configured with the following hardware:

- One virtual CPU (vCPU)
- 2 GB of virtual machine memory
- One virtual network interface card (We have configured this card to use the VMXNET3 driver.)
- One virtual disk for the operating system and RD Web Access role

Figure 9.30 shows the configuration of the RD Web Access virtual machine.

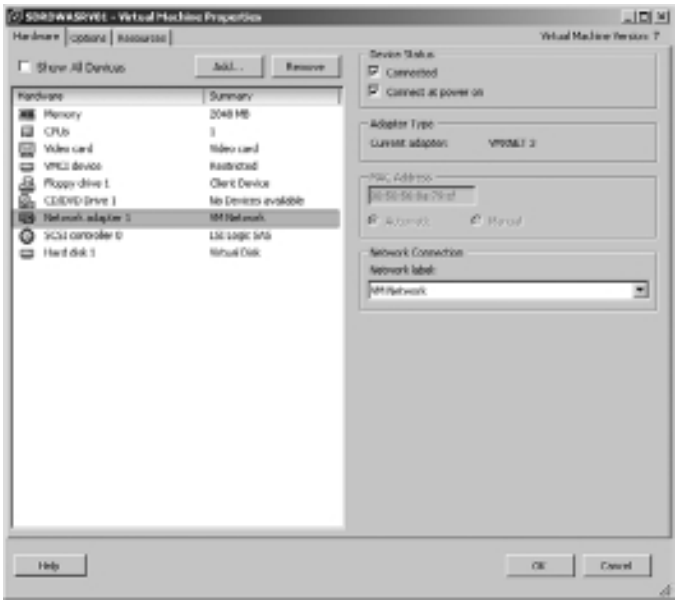

FIGURE 9.30
RD Web Access virtual machine

You install all the required operating service packs and critical updates and join the virtual machine to the Source Discovery domain before adding the RD Web Access server role. Next you use the Add Roles Wizard in Server Manager to install the RD Web Access role service. Follow the steps in the wizard. Make sure you select the Remote Desktop Web Access check box on the Select Role Services page (Figure 9.31).

RD Web Access needs to have requisite service roles and features installed. The installation prompts you to add Web Server (IIS) and the Remote Server Administration Tools, if not already installed (Figure 9.32).

Click Next on the web server (IIS) page (Figure 9.33).

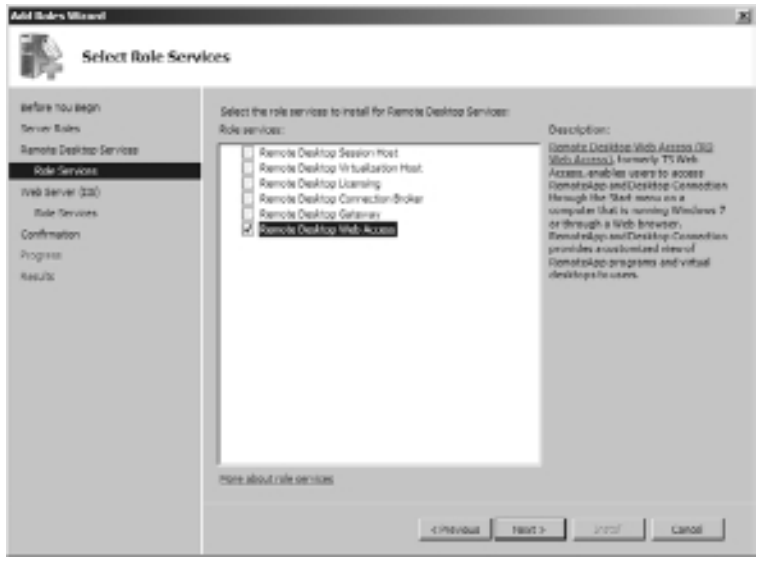

FIGURE 9.31
RD Web Access role service

Add Roles Wizard
Add role services and features required for Remote Desktop Web Access?
You cannot install Remote Desktop Web Access unless the required role services and features are also installed.
Role Services:
Web Server (IIS)
Web Server
Management Tools
Remote Server Administration Tools
Role Administration Tools
Description:
Add Required Role Services
Cancel
Why are these role services and features required?

FIGURE 9.32
RD Web Access prerequisite role services and features

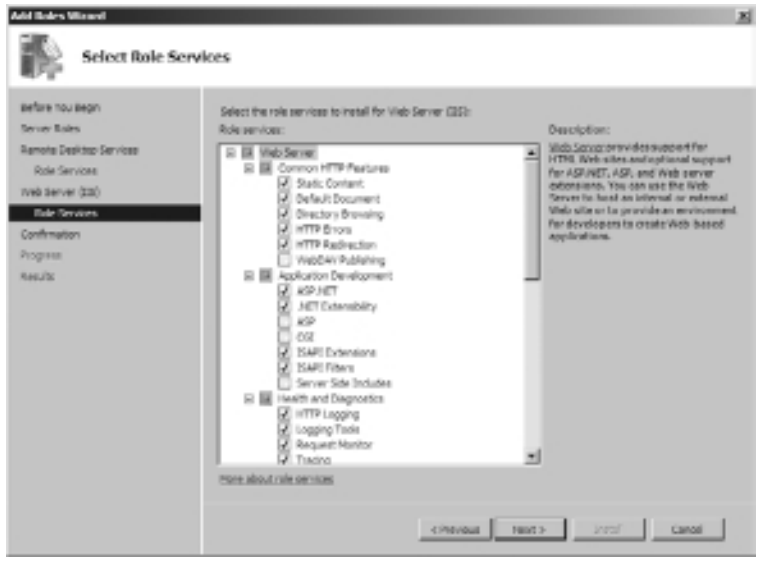

FIGURE 9.33
Web server (IIS) role services

Next confirm your installation selections, and proceed with the installation. This will start the RD Web Access role installation process. The installation progress window displays the role services and features that are installed (Figure 9.34).

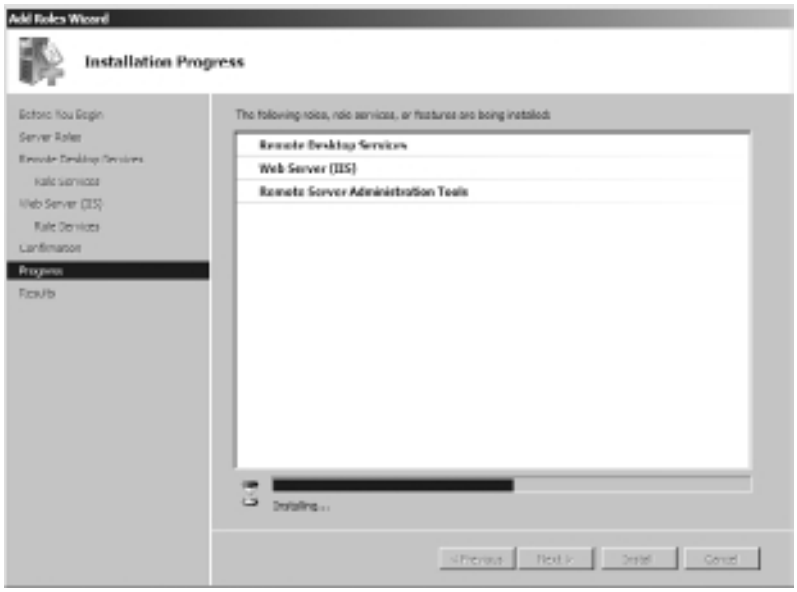

FIGURE 9.34 Installation progress

After installation is complete, make sure that RD Web Access role is successfully installed (Figure 9.35).

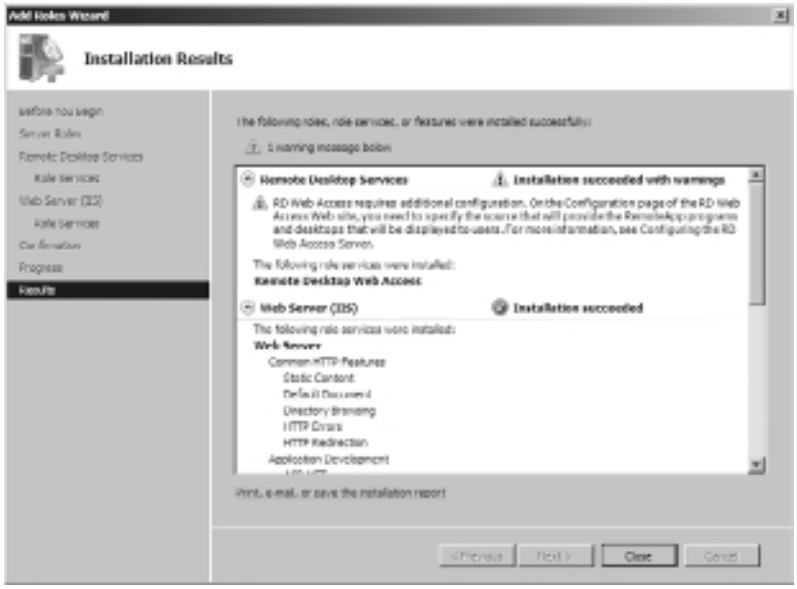
Add Roles Wizard
Installation Results
Before You Begin
Server Roles
Remote Desktop Services
Role Services
Web Server (IIS)
Role Services
Confirmation
Progress
Results
The following roles, role services, or features were installed successfully:
1 warning message below
Remote Desktop Services
Installation succeeded with warnings
RD Web Access requires additional configuration. On the Configuration page of the RD Web Access Web site, you need to specify the source that will provide the RemoteApp programs and desktops that will be displayed to users. For more information, see Configuring the RD Web Access Server.
The following role services were installed:
Remote Desktop Web Access
Web Server (IIS)
Installation succeeded
The following role services were installed:
Web Server
Common HTTP Features
Static Content
Default Document
Directory Browsing
HTTP Errors
HTTP Redirection
Application Development
Print, e-mail, or save the installation report
< Previous
Next >
Close
Cancel

FIGURE 9.35 RD Web Access installation results

Configuring the RD Connection Broker Role

To provide the load-balancing service for the two RD Session Host servers you deployed in the previous section, you need to configure the RD Connection Broker role. You can use the same Windows Server 2008 R2 template as discussed in Chapter 3 to build the RD Connection Broker virtual machine for the proof of concept. Start with the following resources for the virtual machine:

- One virtual CPU (vCPU)
- 2 GB of virtual machine memory
- One virtual network interface card (vNIC); we have configured this card to use the VMXNET3 driver
- One virtual disk for the operating system and RD Connection Broker role

Figure 9.36 shows the configuration of RD Connection Broker virtual machine.

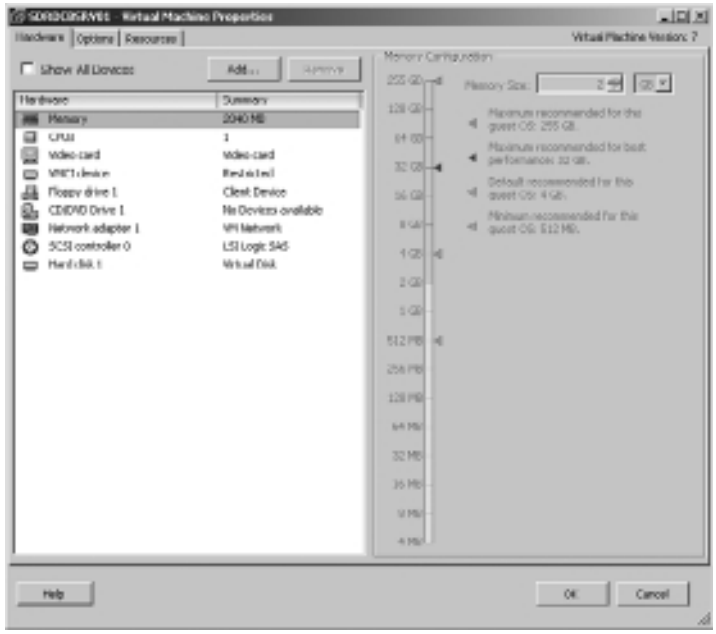

FIGURE 9.36
RD Connection Broker virtual machine configuration

After you install all the required operating service packs and critical updates and join the virtual machine to the Source Discovery domain, you use the Add Roles Wizard in Server Manager to install the RD Connection Broker role service. Follow the steps in the wizard. Make sure you select the Remote Desktop Connection Broker check box on the Select Role Services page (Figure 9.37).

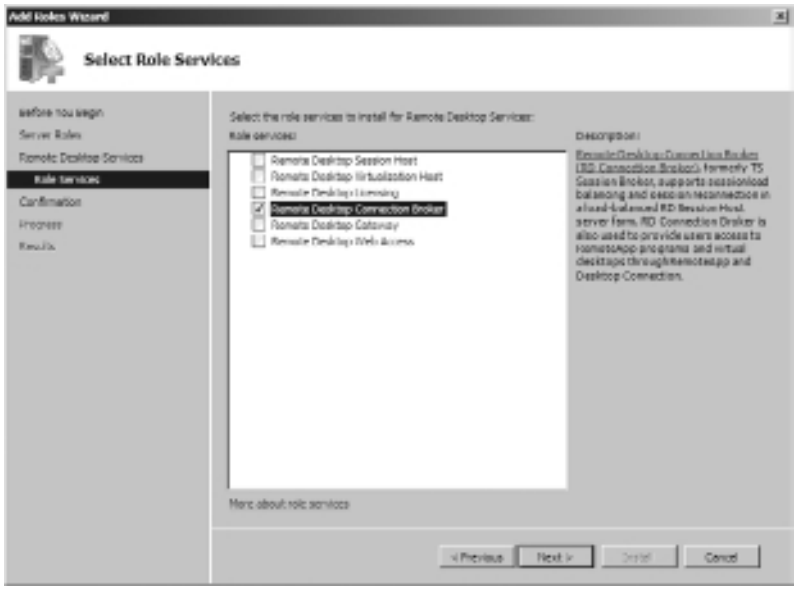
Add Roles Wizard
Select Role Services
Before You Begin
Server Roles
Remote Desktop Services
Role Services
Confirmation
Progress
Results
Select the role services to install for Remote Desktop Services:
Role services:
Remote Desktop Session Host
Remote Desktop Virtualization Host
Remote Desktop Licensing
Remote Desktop Connection Broker
Remote Desktop Gateway
Remote Desktop Web Access
Description:
Remote Desktop Connection Broker (RD Connection Broker), formerly TS Session Broker, supports session load balancing and session reconnection in a load-balanced RD Session Host server farm. RD Connection Broker is also used to provide users access to RemoteApp programs and virtual desktops through RemoteApp and Desktop Connection.
More about role services
< Previous
Next >
Install
Cancel

FIGURE 9.37
RD Connection Broker role service

Next confirm your installation selections, and proceed with the installation (Figure 9.38). This will start the RD Connection Broker role installation process. After installation is complete, click Close.

After installation is complete, make sure that the RD Connection Broker role is successfully installed (Figure 9.39).

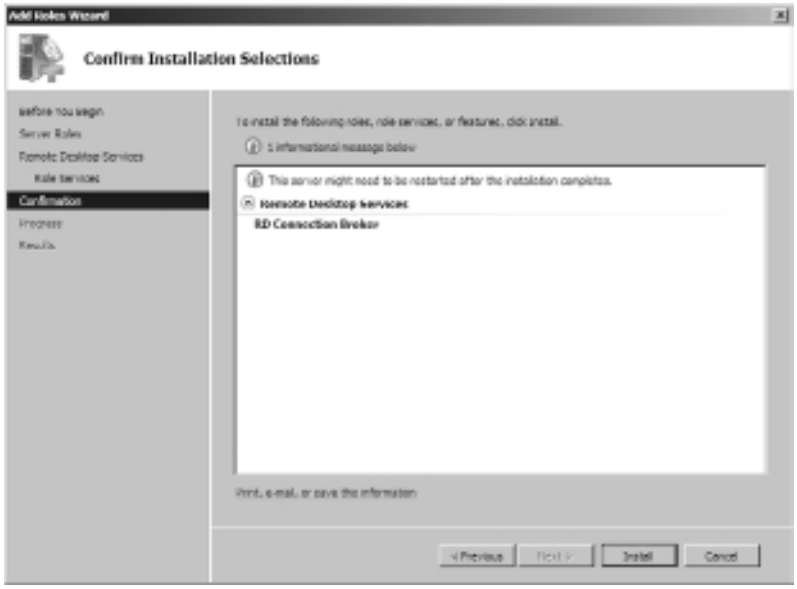

FIGURE 9.38
RD Connection Broker installation selections

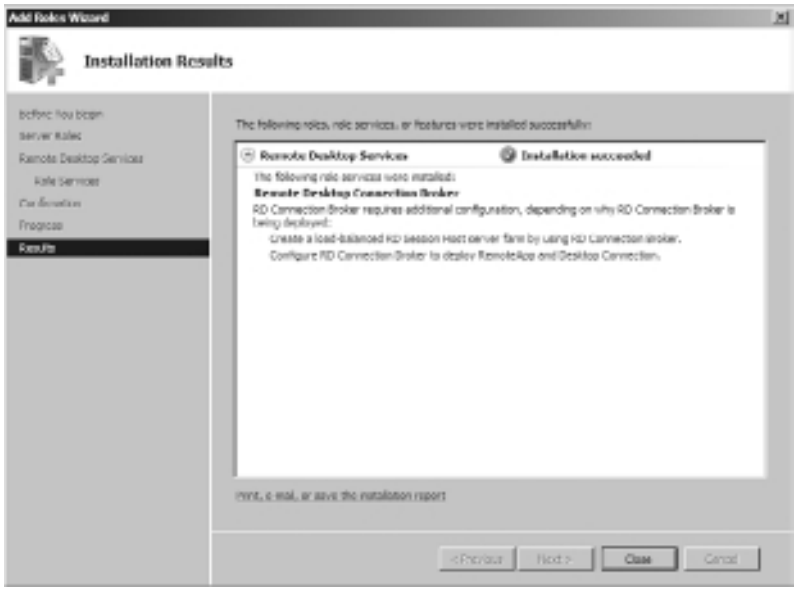
Add Roles Wizard
Installation Results
Before You Begin
Server Roles
Remote Desktop Services
Role Services
Confirmation
Progress
Results
The following roles, role services, or features were installed successfully:
Remote Desktop Services
Installation succeeded
The following role services were installed:
Remote Desktop Connection Broker
RD Connection Broker requires additional configuration, depending on why RD Connection Broker is being deployed:
Create a load-balanced RD Session Host server farm by using RD Connection Broker.
Configure RD Connection Broker to deploy RemoteApp and Desktop Connection.
Print, e-mail, or save the installation report
< Previous
Next >
Close
Cancel

FIGURE 9.39
RD Connection Broker installation results

For a high-availability configuration, you can make sure that vSphere HA is enabled for the vSphere cluster where the RD Connection Broker virtual machines will be run. To provide load balancing and redundancy to the proof-of-concept deployment, you configure an RD Session Host server farm. You start by adding the two RD Session Host servers to the Session Broker Computers group on the RD Connection Broker server (Figure 9.40).

FIGURE 9.40
RD Connection Broker – adding to Session Broker Computers group

Figure 9.41 shows the configuration after these steps.

FIGURE 9.41
RD Connection Broker Session Broker Computers Properties

To create an RD Session Host farm, you need to update the RD Connection Broker settings for each RD Session Host server; you can use the RD Session Host Configuration tool to do this (Figure 9.42).

FIGURE 9.42
RD Connection Broker setting on RD Session Host

On the RD Connection Broker tab, click Change Settings to join the RD Session Host to an RDS farm. You must provide FQDNs for the RD Connection Broker and the RDS farm (Figure 9.43).

RD Connection Broker Settings
Select how you want this Remote Desktop Session Host server to be used with RD Connection Broker.
Remote Desktop Virtualization
Virtual machine redirection
Provides redirection for virtual machines used in RemoteApp and Desktop Connection.
Remote Desktop Services
Dedicated farm redirection
Provides dedicated redirection for the specified farm.
Farm member
Joins this Remote Desktop Session Host server to the specified farm.
No farm membership or redirection
RD Connection Broker server name:
Farm Name:
More about configuring redirection and farm membership
OK
Cancel

FIGURE 9.43
RD Session Host farm configuration

Enable Connection Broker Load Balancing, and select the RD Session Host IP address for reconnection (Figure 9.44).

FIGURE 9.44
RD Connection Broker Load Balancing

You need to repeat these RD Connection Broker settings on each RD Session Host in your RDS farm (Figure 9.45).

FIGURE 9.45
RD Session Host farm configuration

You need to add the RDS farm name to Active Directory to enable DNS round-robin load balancing for new user connections. Figure 9.46 shows the Active Directory DNS Manager after these steps are completed.

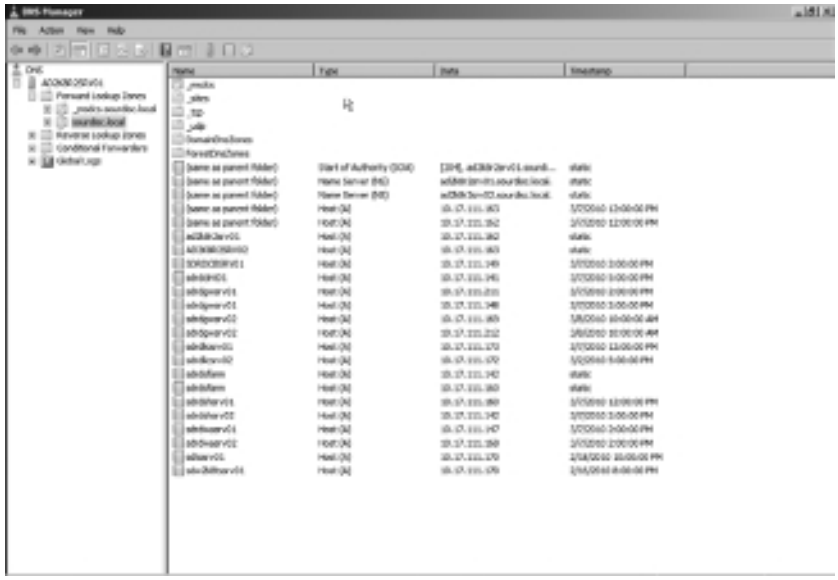

FIGURE 9.46
RDS farm DNS round-robin configuration

Configuring RemoteApp

RemoteApp functionality allows end users to access only specific applications from their web browser. Configuring a RemoteApp program is fairly straightforward process. You need to have appropriate privileges in place before end users can access these remote applications. In this section, we will first set these privileges and then configure RemoteApp for Microsoft Word, PowerPoint, and Excel applications.

You need to add the RD Web Access computer account object to the TS Web Access Computers security group on the RD Connection Broker server (Figure 9.47).

FIGURE 9.47
RD Connection Broker – adding to Web Access Computers group

Figure 9.48 shows the configuration after these steps.

FIGURE 9.48
RD Connection Broker Web Access Computers Properties

You also need to add the RD Connection Broker server to the TS Web Access Group on the two RD Session Host servers (Figure 9.49).

FIGURE 9.49
RD Session Host – adding to Web Access Computers group

Figure 9.50 shows the configuration after these steps.

TS Web Access Computers Properties
General
TS Web Access Computers
Description:
Members of this group can query the list of RemoteApp programs on this computer, including those available
Members:
Add...
Remove
Changes to a user's group membership are not effective until the next time the user logs on.
OK
Cancel
Apply
Help

FIGURE 9.50
RD Session Host Web Access Computers Properties group

After adding the required privileges, you configure the RemoteApp programs on RD Session Host using RemoteApp Manager (Figure 9.51).

FIGURE 9.51
RemoteApp Manager

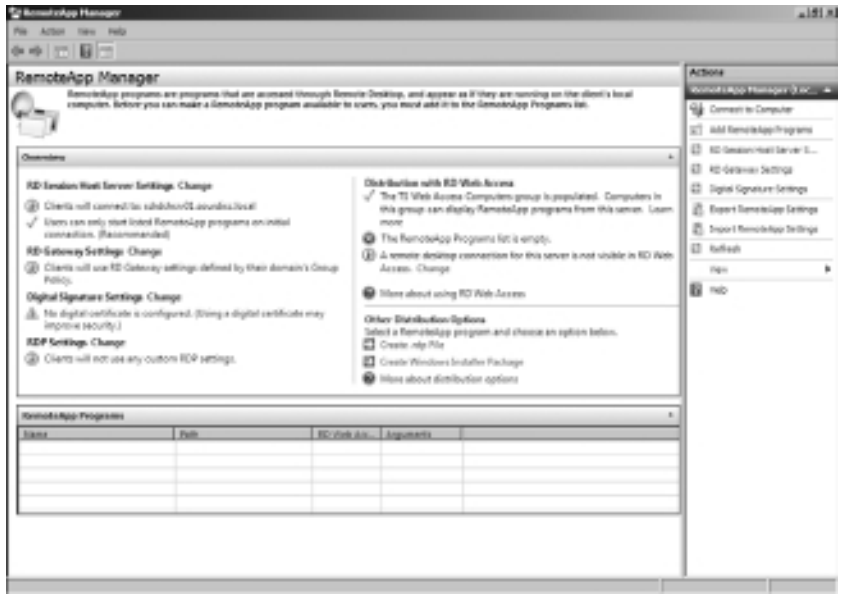

Click Add RemoteApp Programs in the Actions pane to start the Welcome to the RemoteApp Wizard page (Figure 9.52).

FIGURE 9.52
Welcome To The RemoteApp Wizard page

You add Microsoft Office Excel 2007, Microsoft Office PowerPoint 2007, and Microsoft Office Word 2007 applications to the RemoteApp programs list (Figure 9.53).

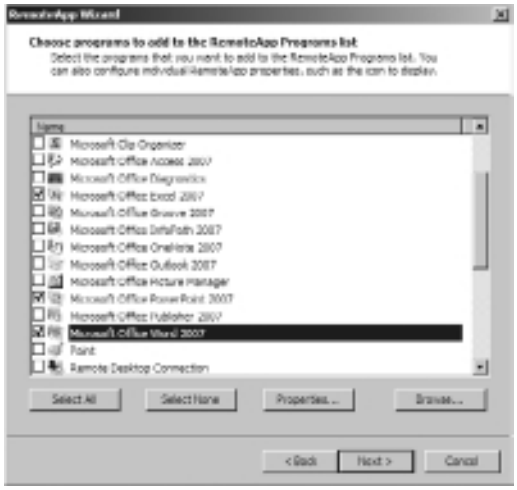
RemoteApp Wizard
Choose programs to add to the RemoteApp Programs list
Select the programs that you want to add to the RemoteApp Programs list. You can also configure individual RemoteApp properties, such as the icon to display.
Name
Microsoft Clip Organizer
Microsoft Office Access 2007
Microsoft Office Diagnostics
Microsoft Office Excel 2007
Microsoft Office Groove 2007
Microsoft Office InfoPath 2007
Microsoft Office OneNote 2007
Microsoft Office Outlook 2007
Microsoft Office Picture Manager
Microsoft Office PowerPoint 2007
Microsoft Office Publisher 2007
Microsoft Office Word 2007
Paint
Remote Desktop Connection
Select All
Select None
Properties...
Browse...
< Back
Next >
Cancel

FIGURE 9.53
RemoteApp programs list

You can review the settings on the next page and click Finish to add the programs (Figure 9.54).

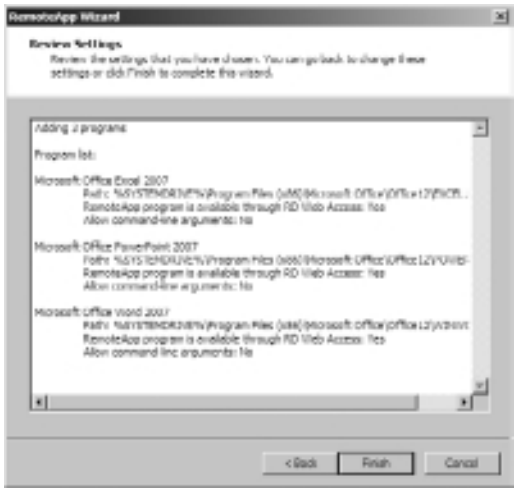
RemoteApp Wizard
Review Settings
Review the settings that you have chosen. You can go back to change these settings or click Finish to complete the wizard.
Adding 3 programs
Program list:
Microsoft Office Excel 2007
Path: %SYSTEMDRIVE%\Program Files (x86)\Microsoft Office\Office12\EXCEL.EXE
RemoteApp program is available through RD Web Access: Yes
Allow command line arguments: No
Microsoft Office PowerPoint 2007
Path: %SYSTEMDRIVE%\Program Files (x86)\Microsoft Office\Office12\POWERPNT.EXE
RemoteApp program is available through RD Web Access: Yes
Allow command line arguments: No
Microsoft Office Word 2007
Path: %SYSTEMDRIVE%\Program Files (x86)\Microsoft Office\Office12\WINWORD.EXE
RemoteApp program is available through RD Web Access: Yes
Allow command line arguments: No
< Back
Finish
Cancel

FIGURE 9.54
RemoteApp Review Settings

RemoteApp Manager lists the three Microsoft Office applications that were added as RemoteApp programs (Figure 9.55).

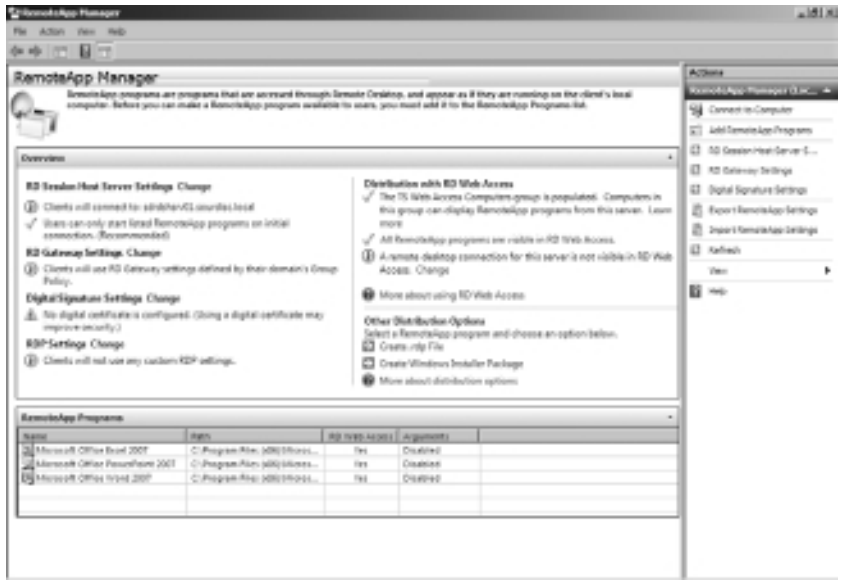

FIGURE 9.55
RemoteApp programs

Next, you assign a single RemoteApp source on the RD Web Access server for these RemoteApp programs using the Remote Desktop Web Access Configuration tool (Figure 9.56). For the proof of concept, you're using a self-signed certificate for the RD Web Access server, so you need to create an exception for a certificate error.

FIGURE 9.56
Remote Desktop Web Access configuration login

On the Configuration page, select an RD Connection Broker server as the source, provide the name, and click OK (Figure 9.57).

FIGURE 9.57
Remote Desktop Web Access configuration

Finally, you add a RemoteApp source on the RD Connection Broker by using Remote Desktop Connection Manager (Figure 9.58).

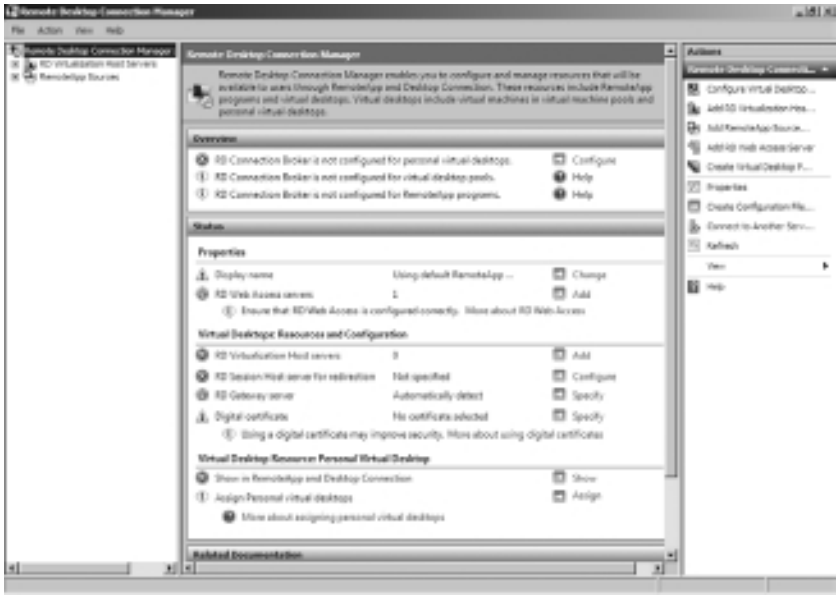

FIGURE 9.58
Remote Desktop Connection Manager

Because you are using an RD Session Host server farm, you add the name of the RDS farm as the RemoteApp source name (Figure 9.59).

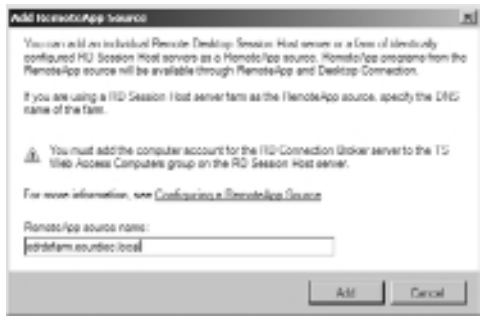

FIGURE 9.59
Remote Desktop Connection Manager – adding RemoteApp source

Figure 9.60 shows the completed configuration.

FIGURE 9.60
Remote Desktop Connection Manager – list of RemoteApp Sources

Configuring the RD Gateway Role

You need to deploy the RD Gateway role to enable external users to connect to corporate resources inside the firewall without the need to log into the VPN connection. To deploy the RD Gateway role, once again you will leverage the Windows Server 2008 R2 template you created in Chapter 3.

For high availability and fault tolerance, you will configure two RD Gateway virtual machines using Windows Network Load Balancing (NLB). The RD Web Access virtual machine is configured with the following hardware:

- One virtual CPU (vCPU)
- 2 GB of virtual machine memory
- Two virtual network interface cards: one public network interface card for client access and cluster management and one heartbeat network interface card for cluster communications
- One virtual disk for the operating system and RD Gateway role

Figure 9.61 shows the configuration of the RD Gateway virtual machine.

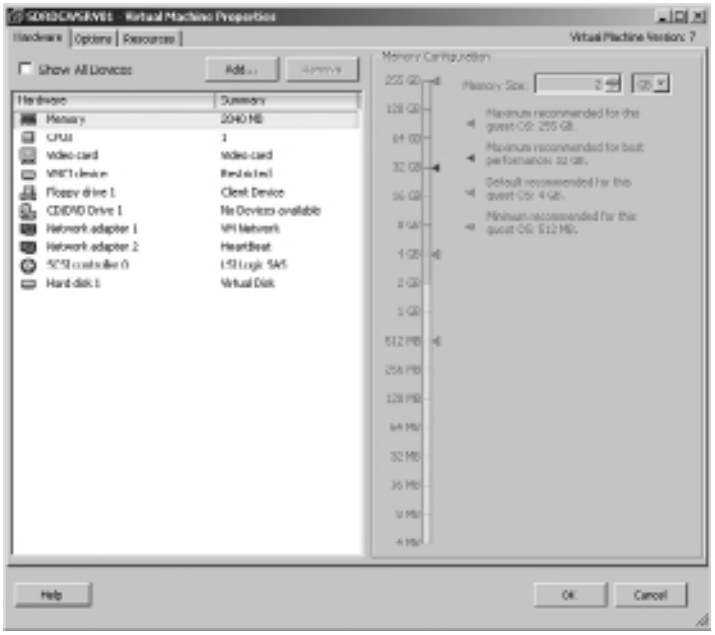

FIGURE 9.61
RD Gateway server role virtual machine configuration

You install all the required operating service packs and critical updates and join the virtual machine to the Source Discovery domain before adding the RD Web Access server role. Next you use the Add Roles Wizard in Server Manager to install the RD Gateway role service. Follow the steps in the wizard. Make sure you select the Remote Desktop Web Gateway check box on the Select Role Services page (Figure 9.62).

When installing the RD Gateway role service, Server Manager will automatically prompt for and install additional roles, role services, and features that are not already installed (Figure 9.63):

- Remote procedure call (RPC) over HTTP Proxy
- Web Server (IIS)
- Network Policy and Access Services

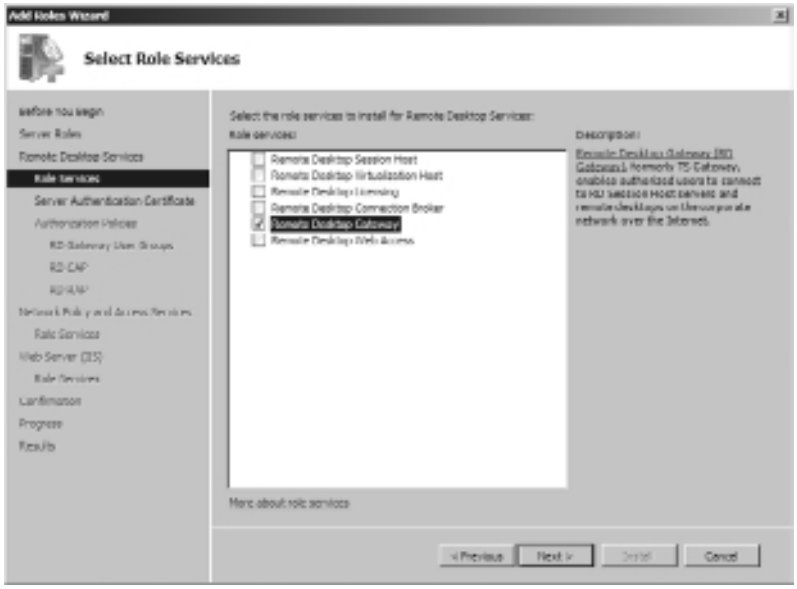

FIGURE 9.62
RD Gateway role service

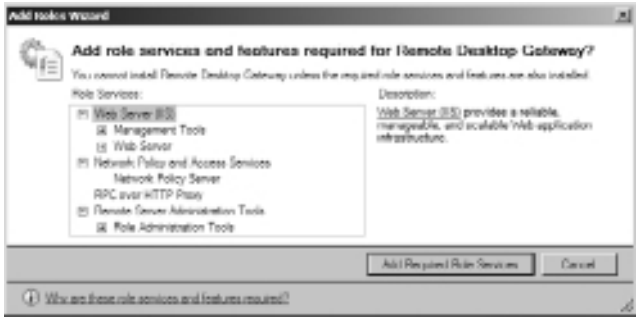
Add Roles Wizard
Add role services and features required for Remote Desktop Gateway?
Role Services:
Web Server (IIS)
Management Tools
Web Server
Network Policy and Access Services
Network Policy Server
RPC over HTTP Proxy
Remote Server Administration Tools
Role Administration Tools
Description:
Web Server (IIS) provides a reliable, manageable, and scalable Web application infrastructure.
Add Required Role Services
Cancel
Why are these role services and features required?

FIGURE 9.63
RD Web Access prerequisite role services and features

For the proof of concept, you choose to create a self-signed certificate for SSL encryption (Figure 9.64).

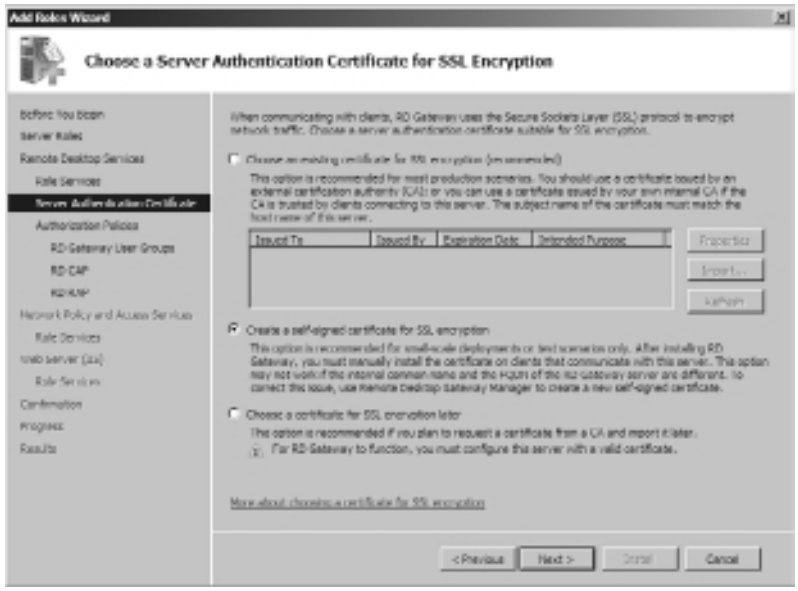
Add Roles Wizard
Choose a Server Authentication Certificate for SSL Encryption
Before You Begin
Server Roles
Remote Desktop Services
Role Services
Server Authentication Certificate
Authorization Policies
RD Gateway User Groups
RD CAP
RD RAP
Network Policy and Access Services
Role Services
Web Server (IIS)
Role Services
Confirmation
Progress
Results
When communicating with clients, RD Gateway uses the Secure Sockets Layer (SSL) protocol to encrypt network traffic. Choose a server authentication certificate suitable for SSL encryption.
Choose an existing certificate for SSL encryption (recommended)
This option is recommended for most production scenarios. You should use a certificate issued by an external certification authority (CA); or you can use a certificate issued by your own internal CA if the CA is trusted by clients connecting to this server. The subject name of the certificate must match the host name of this server.
Issued To
Issued By
Expiration Date
Intended Purpose
Properties
Import...
Refresh
Create a self-signed certificate for SSL encryption
This option is recommended for small-scale deployments or test scenarios only. After installing RD Gateway, you must manually install the certificate on clients that communicate with this server. This option may not work if the internal common name and the FQDN of the RD Gateway server are different. To correct this issue, use Remote Desktop Gateway Manager to create a new self-signed certificate.
Choose a certificate for SSL encryption later
This option is recommended if you plan to request a certificate from a CA and import it later.
For RD Gateway to function, you must configure this server with a valid certificate.
More about choosing a certificate for SSL encryption
< Previous
Next >
Install
Cancel

FIGURE 9.64
Server authentication certificate

On the Create Authorization Policies For RD Gateway page, select Now, and then click Next (Figure 9.65).

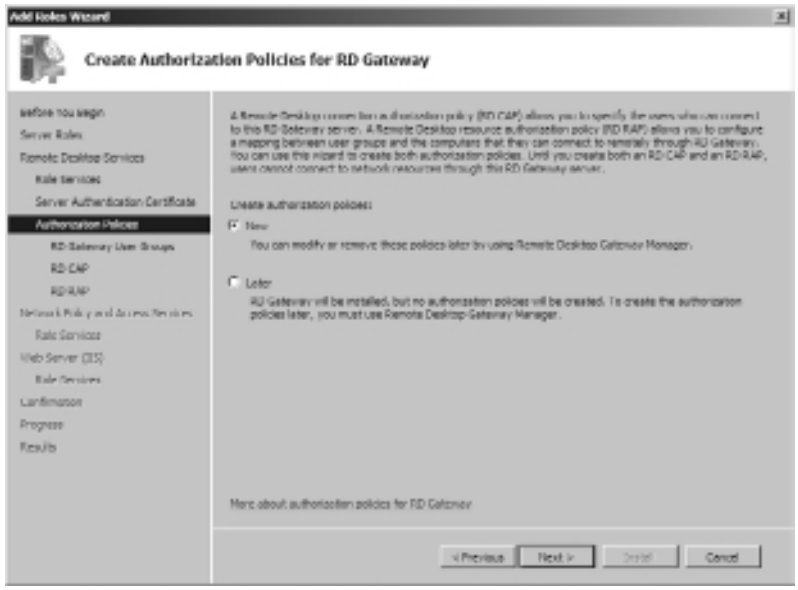

FIGURE 9.65
Authorization policies

You select all domain users to be part of the RD Gateway User groups (Figure 9.66).

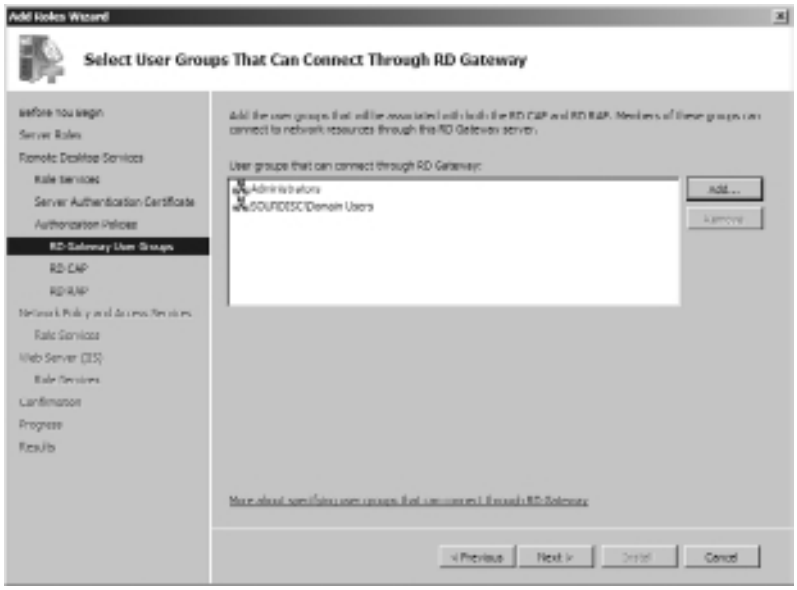

FIGURE 9.66
RD Gateway User groups

You create an RD connection authorization policy for RD Gateway using the password authentication method (Figure 9.67).

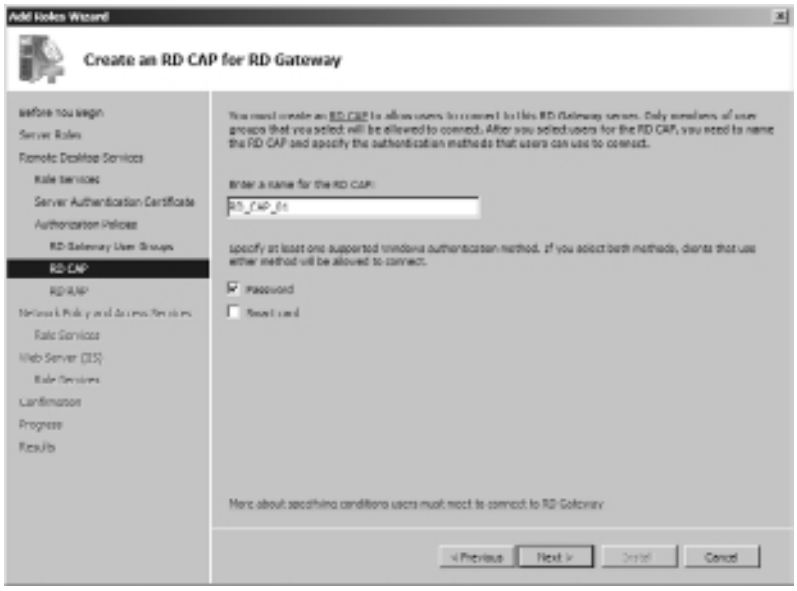

FIGURE 9.67
RD CAP for RD Gateway

You create an RD resource authorization policy for RD Gateway for all users to connect to any computer on the network (Figure 9.68).

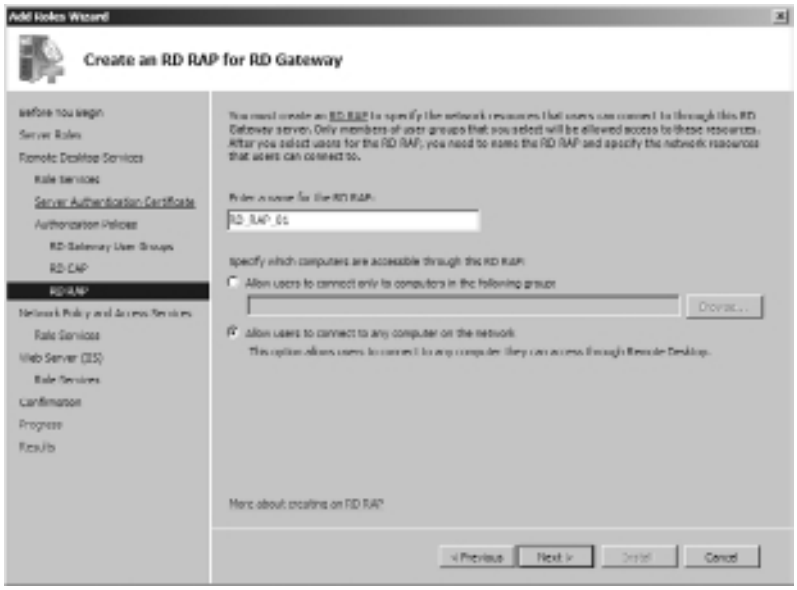

FIGURE 9.68
RD RAP for RD Gateway

Verify that Network Policy Server is selected, and then click Next (Figure 9.69).

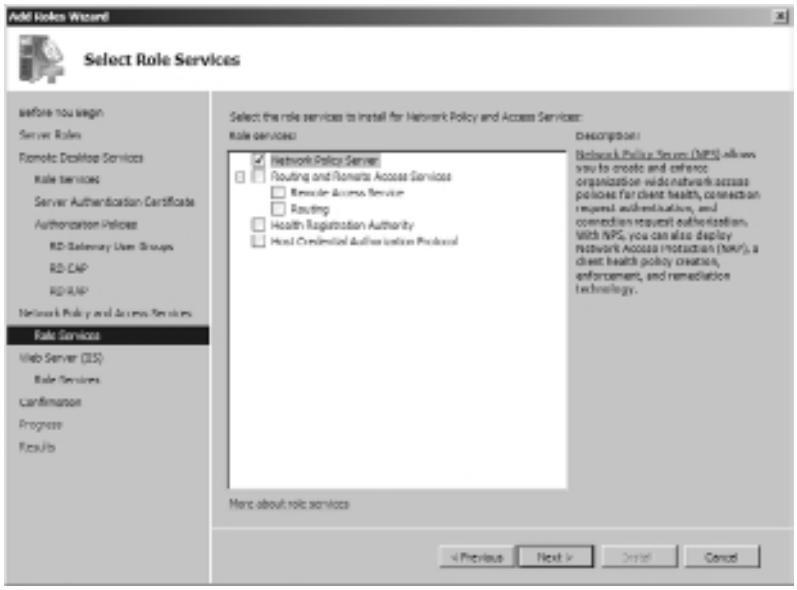
Add Roles Wizard
Select Role Services
Before You Begin
Server Roles
Remote Desktop Services
Role Services
Server Authentication Certificate
Authorization Policies
RD Gateway User Groups
RD CAP
RD RAP
Network Policy and Access Services
Role Services
Web Server (IIS)
Role Services
Confirmation
Progress
Results
Select the role services to install for Network Policy and Access Services:
Role services:
Network Policy Server
Routing and Remote Access Services
Remote Access Service
Routing
Health Registration Authority
Host Credential Authorization Protocol
Description:
Network Policy Server (NPS) allows you to create and enforce organization-wide network access policies for client health, connection request authentication, and connection request authorization. With NPS, you can also deploy Network Access Protection (NAP), a client health policy creation, enforcement, and remediation technology.
More about role services
< Previous
Next >
Install
Cancel

FIGURE 9.69
Network Policy Server Role service

Click Next on the web server (IIS) page (Figure 9.70).

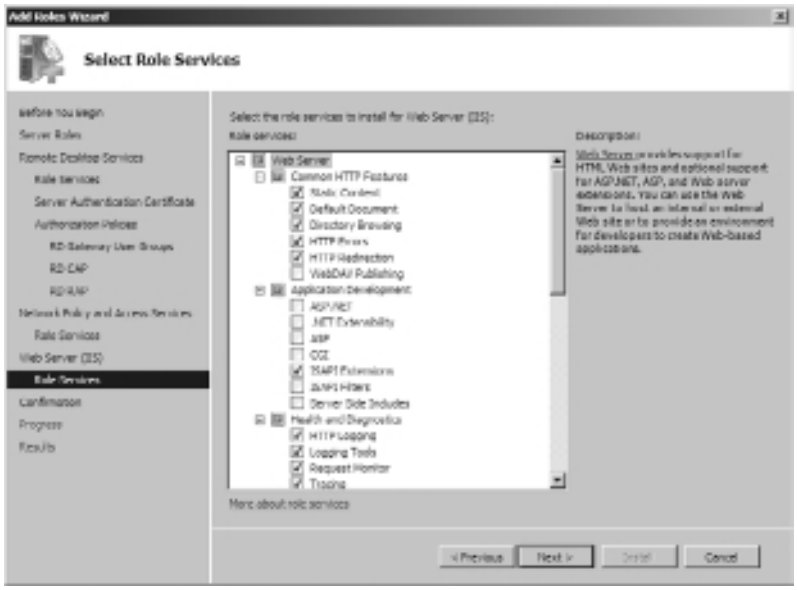

Next confirm your installation selections, and proceed with the installation. The installation progress window displays the role services and features that are installed (Figure 9.71).

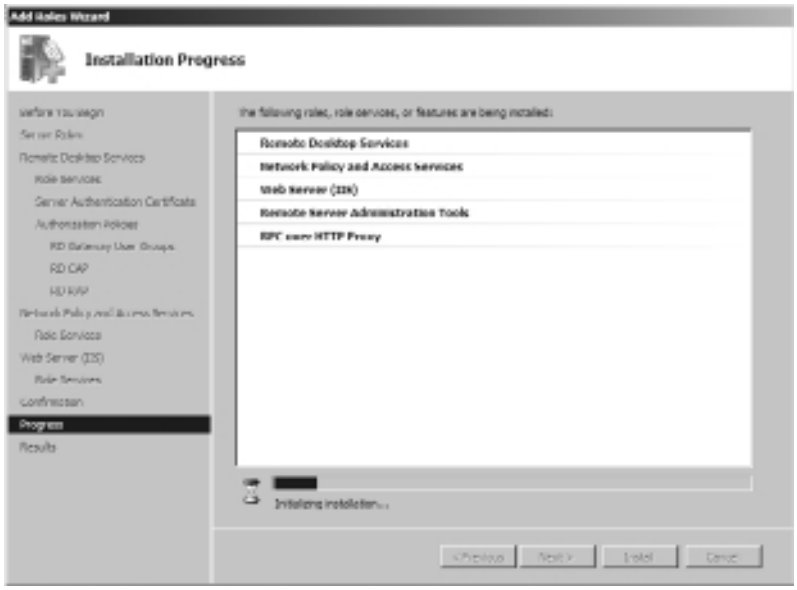

After installation is complete, make sure that all the roles and features are successfully installed (Figure 9.72).

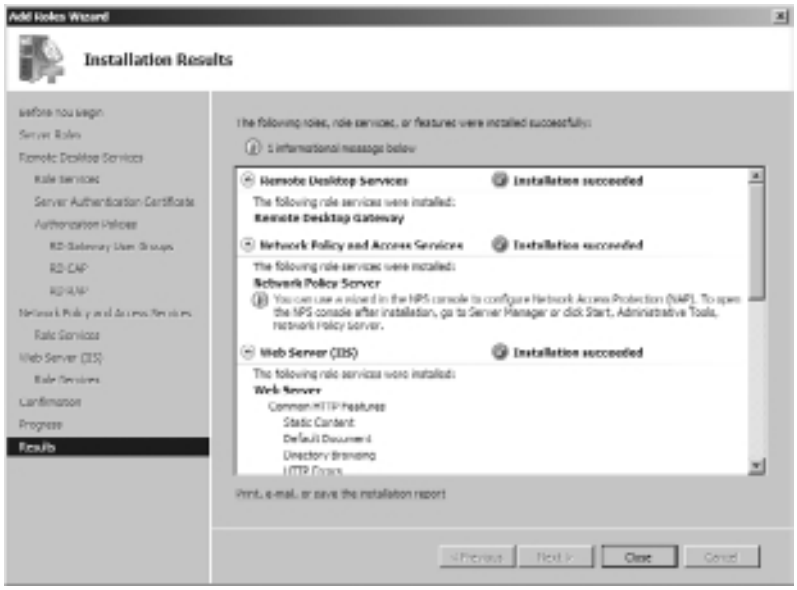

FIGURE 9.72
RD Gateway installation results

Configuring Windows Network Load Balancing

End users access an RDS farm via the RD Gateway and RD Web Access servers, so these services need to be highly available. You can utilize Windows Network Load Balancing (WNLB), third-party software load-balancing solutions, or hardware load balancers. Using WNLB is a good alternative to using a hardware load-balancing solution. For smaller RDS farms and development and test environments, this can provide excellent cost savings. The downside to using WNLB is that the load balancing is handled by the networking stack, which in large RDS farms can introduce some performance degradation. In this section, we will show how to deploy WNLB for RD Gateway servers. We use the second provisioned virtual NIC on each RD Gateway virtual machine to create this WNLB-based load-balancing configuration. You can follow similar steps to configure WNLB for RD Web Access servers.

TIP For larger RDS farm deployments where there are more than two RD Gateway roles, we recommend using a hardware load-balancing solution because all the load-balancing operations will occur on the hardware and not the virtual machine.

You must first enable the Network Load Balancing feature on both RD Gateway servers using Server Manager (Figure 9.73).

After installation is complete, make sure that Network Load Balancing feature is successfully installed (Figure 9.74).

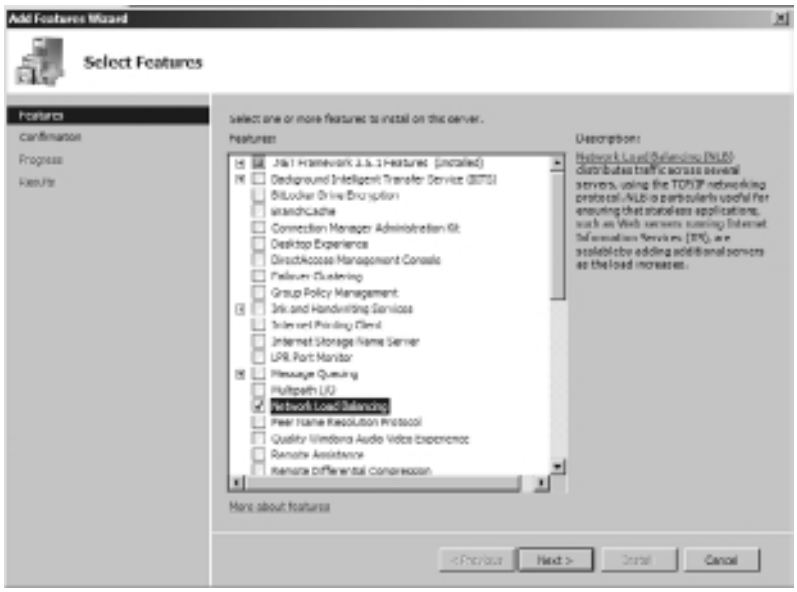

FIGURE 9.73
Adding the Network Load Balancing feature

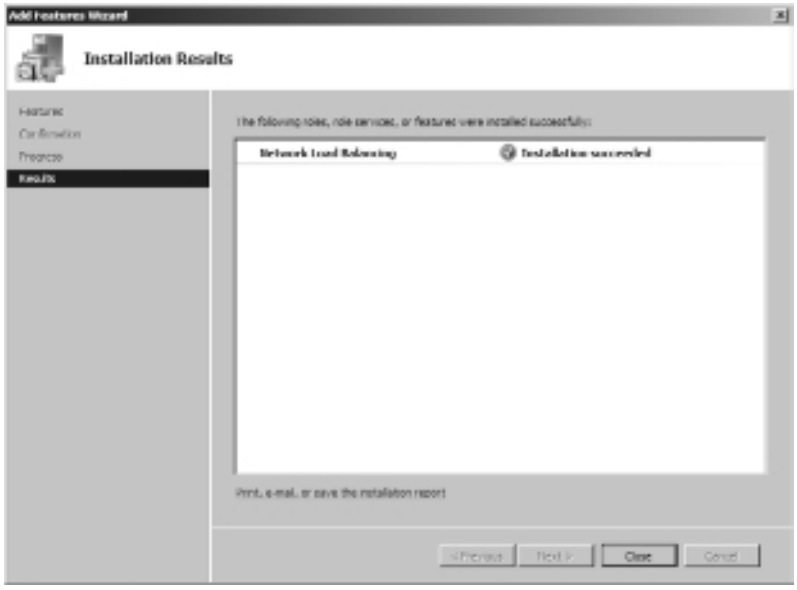

FIGURE 9.74
Network Load Balancing installation results

To create a new load-balancing cluster, launch Network Load Balancing Manager from Administration Tools (Figure 9.75).

FIGURE 9.75
New network load-balancing cluster

In the Host field, place the name of the first RD Gateway server host, and click the Connect button to connect to and configure the first host in the Windows load-balancing cluster (Figure 9.76).

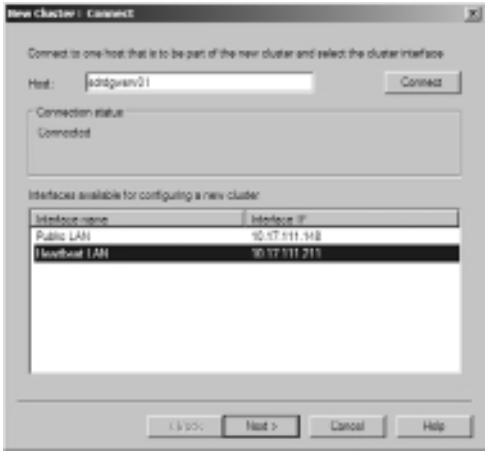

FIGURE 9.76
Entering the first RD Gateway server to create the WNLB cluster

Choose the second virtual NIC as the WNLB interface that will be used for the load balancing (Figure 9.77). The IP address for the second virtual NIC was entered ahead of time after the configuration of the Windows operating system.

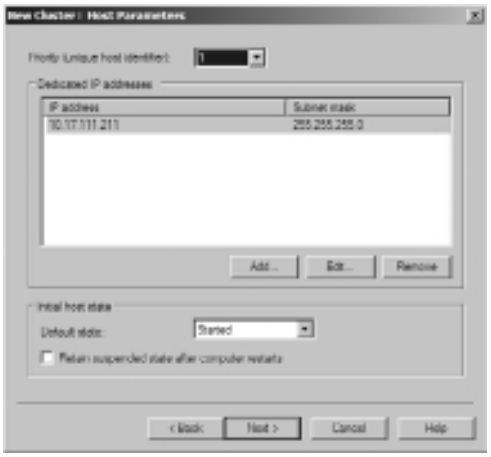

FIGURE 9.77
First RD Gateway server parameters for the WNLB cluster

WARNING When configuring the NLB virtual NIC, do not select the Network Load Balancing (NLB) check box, because doing so will prevent the Windows Load Balancing Manager from being able to access the virtual NIC for configuration.

After clicking the Next box to continue the installation, you can then enter the IP address and subnet mask for the cluster (Figure 9.78).

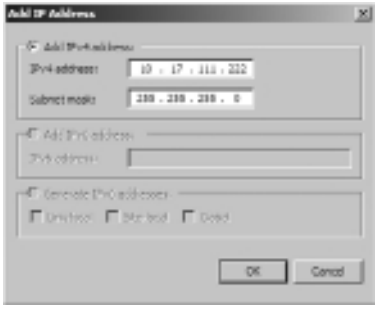

FIGURE 9.78
Choosing the cluster IP address

Clicking the OK button takes you to the next screen (Figure 9.79) where you must enter a fully qualified DNS name for your WNLB cluster. After entering the fully qualified DNS name for the cluster, you select the Unicast option for the cluster operation mode. This will enable the WNLB cluster to utilize the second provisioned virtual NIC.

FIGURE 9.79
Cluster IP parameters

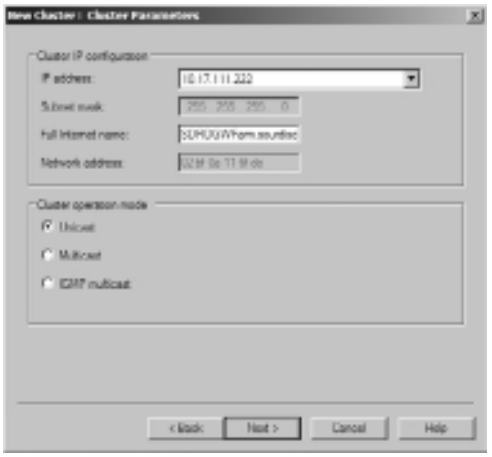

On the next screen you have the option of specifying which traffic will go through the WNLB cluster (Figure 9.80). Edit these rules to select the TCP protocol, and leave the others at the defaults. You continue by clicking the Finish button.

FIGURE 9.80
Cluster port rules

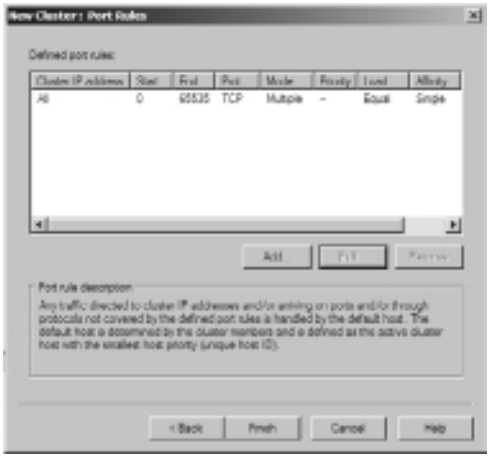

At this point you have created the Windows NLB cluster with a single RD Gateway server. In this configuration, you have a cluster but a single point of failure, with no load balancing or redundancy for the RD Gateway server role. If you were to experience an RD Gateway server failure, external Source Discovery end users would not be able to access the RDS farm unless they used VPN to connect to the corporate network. This would not be an acceptable solution and would defeat the purpose of deploying WNLB in the first place. So, to continue, you now add the second web RD Gateway server to the existing WNLB cluster configuration by selecting Add Host To Cluster from the Cluster menu (Figure 9.81).

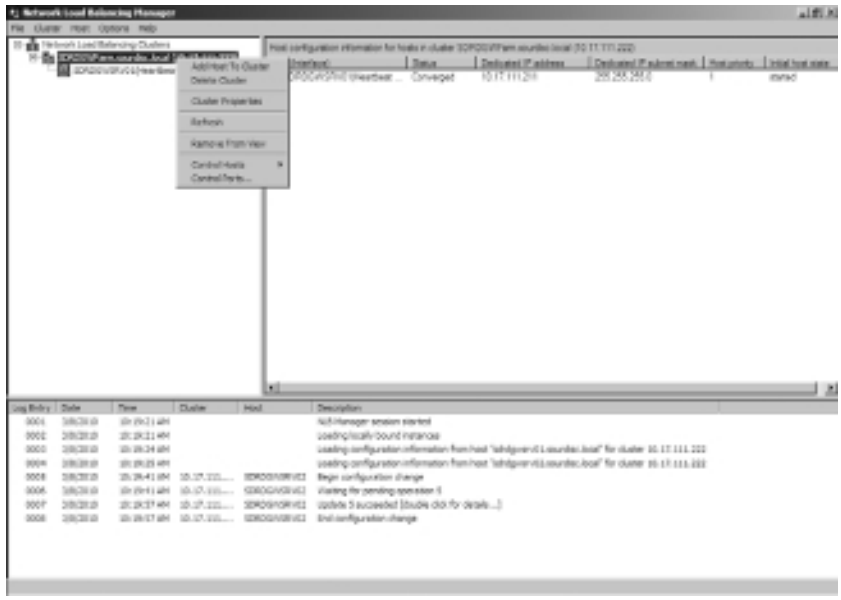

FIGURE 9.81
Adding the second RD Gateway server to the Windows NLB cluster

As you can see, the process is the same as adding the first RD Gateway server to form the cluster. In the Host field, you place the name of the second RD Gateway server host, and click the Connect button to connect to and configure the second host in the Windows load-balancing cluster (Figure 9.82).

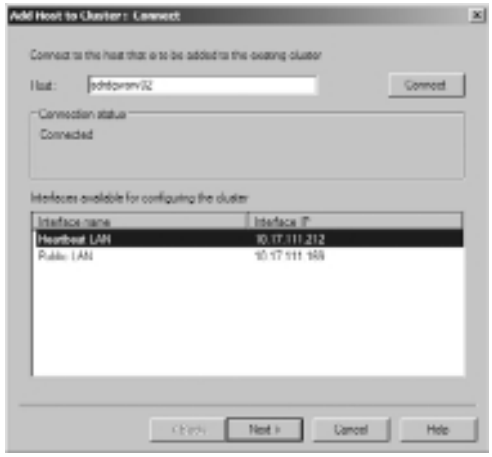

FIGURE 9.82
Connecting to the second RD Gateway server to create the WNLB cluster

Make sure the second virtual NIC is used as the WNLB interface that will be used for the load balancing (Figure 9.83).

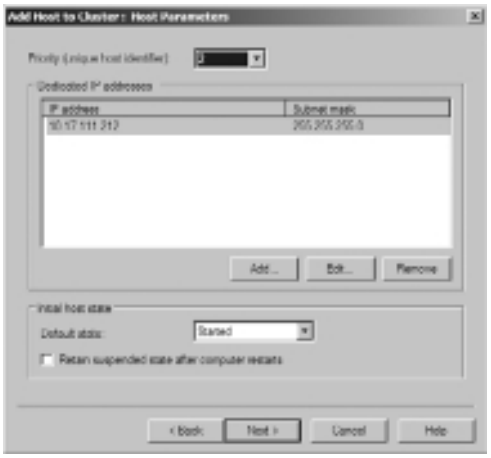

FIGURE 9.83
Second RD Gateway server parameters for the WNLB cluster

Select the default values on the Port Rules screen. Continue by clicking the Finish button (Figure 9.84).

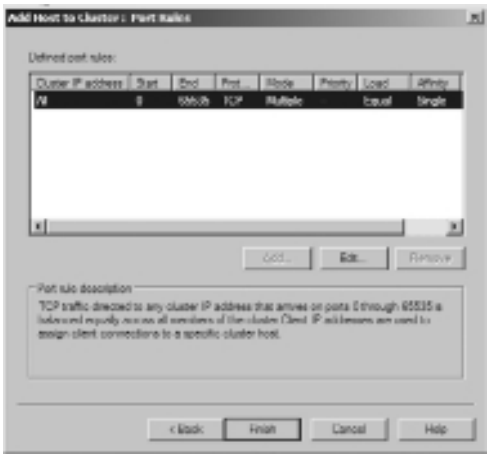

FIGURE 9.84
Add Host to Cluster: Port Rules page

After adding the second host, the cluster will go through a cluster transition state known as convergence in which both cluster nodes are synchronized with one another. It could take up to a minute or more to complete convergence. Once complete, all the nodes of the WNLB cluster should be fully operational and display a green status. Figure 9.85 shows the Windows NLB cluster state after convergence has completed.

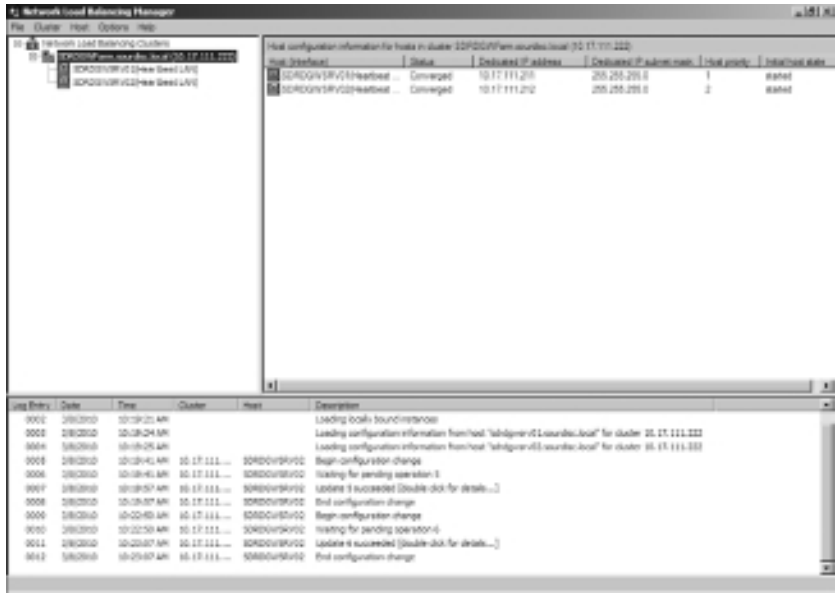

FIGURE 9.85
A fully operational Windows NLB cluster

Next you need to use RD Gateway Manager to add each RD Gateway server to the server farm. You need to perform these steps on each RD Gateway server and make sure that the farm status is OK for each server (Figure 9.86).

FIGURE 9.86
RD Gateway server farm status for first server

Figure 9.87 shows the server farm steps on the second RD Gateway server.

FIGURE 9.87
RD Gateway server farm status for second server

This completes the deployment of the Source Discovery RDS farm that includes two RD Licensing servers, a two-host RD Session Host server farm load-balanced using RD Connection Broker, an RD Gateway Server farm load-balanced using WNLB and an RD Web Access server. In the next section, we will show how to test this virtual RDS farm deployment.

Testing Virtual RDS Farm Deployment

Before you can test the virtual RDS deployment, you export the SSL certificates for the RD Web Access server and RD Gateway server. You copy them to a Windows XP client machine, since most Source Discovery employees still use Windows XP as the client operating system. You then import these SSL certificates on the client machine. You connect to the RemoteApp program using Internet Explorer (Figure 9.88). You also have other options to create and distribute client configuration files to end users.

All the three Microsoft Office 2007 applications are available as RemoteApp programs (Figure 9.89).

FIGURE 9.88
End user connecting to RemoteApp

FIGURE 9.89
RemoteApp programs

Start by launching the Microsoft Word 2007 application (Figure 9.90).

FIGURE 9.90
Running Microsoft Word RemoteApp

The Word application launches successfully and opens a new Word document (Figure 9.91).

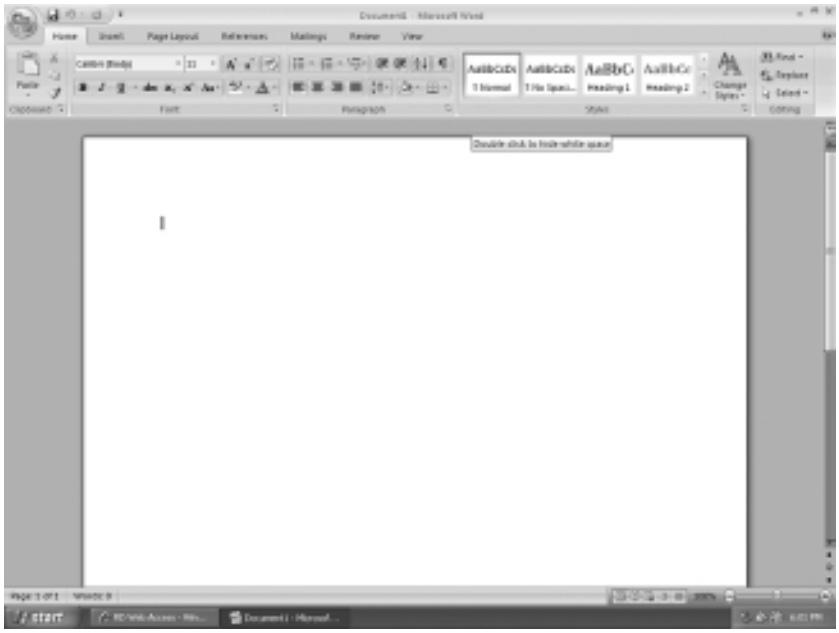

FIGURE 9.91
New Microsoft Word document opened by RemoteApp

Launch the other two RemoteApp programs—Microsoft Excel 2007 and Microsoft PowerPoint 2007 (Figure 9.92).

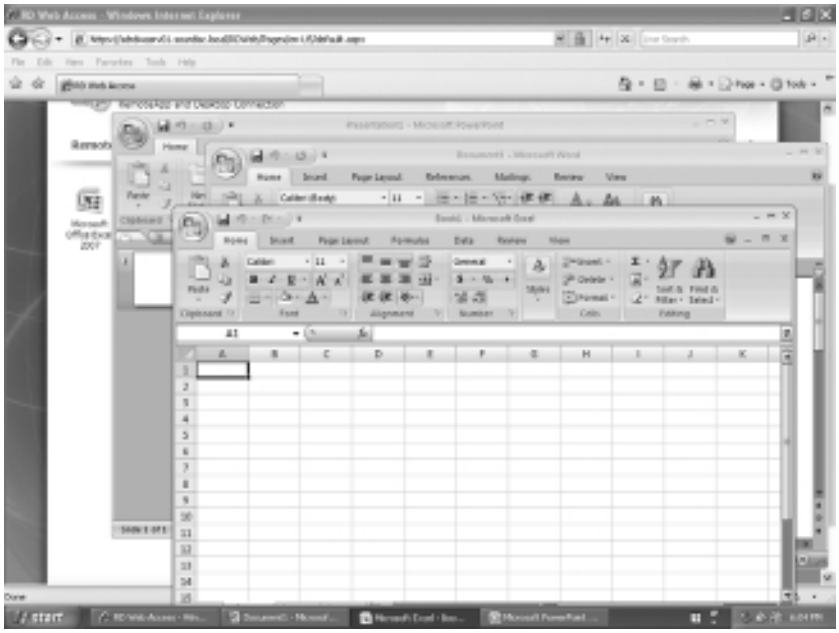

FIGURE 9.92 All three RemoteApp Programs launch successfully.

After testing basic single-user access, you can test the functionality of the RD Session Host farm and RD Gateway server farm. To do this, power off one of the servers in each farm and ensure you can access Source Discovery's RD farm.

Testing RDS deployment using a realistic workload is particularly important, unlike any other data-center enterprise server applications. Plan enough time for the testing. Once your test environment is created, simulate all production environment activities, and monitor the environment. Test current production environment procedures and policies, and adjust and document changes where necessary to support the proposed consolidated environment. You should not only scale to the user count you plan to support; you should do so while replicating exact user actions performed on the applications that you plan to deploy.

Several testing tools, free and commercial, are available to test your deployment. You can start with tools from Microsoft such as Remote Desktop Load Simulation Tools that can assist you with capacity planning for RDS deployment. These tools include automation framework and application scripting support that enables you to simulate user interaction with RDS. This can help you decide whether your RDS environment can handle the load that your organization expects to place on it.

TIP You can download the Remote Desktop Load Simulation Tools from the Microsoft Download Center (`http://go.microsoft.com/fwlink/?LinkId=178956`).

Several other tools are also available to help you with capacity planning exercises for your RDS environment. Customers often record their application workload using HP's LoadRunner tool or scripting products such as AutoIT and replay these scripts to simulate the required user workload.

Like any benchmark products, these performance tools will only help extrapolate your application's performance. Use realistic workloads for the test environment. Stress testing the deployments on VMware vSphere with no "think times," for example, or comparing physical to virtual performance from different physical servers can be misleading. Evaluate the proposed workloads to gain an understanding of how current physical resource utilization peaks and heterogeneous workloads will affect resource utilization when these workloads are virtualized. Add more workloads as appropriate, keeping in mind CPU, memory, network, storage, and other critical RDS-related metrics to ensure that the ESX server or its resources are not over-subscribed.

Test the load-balancing and high-availability configuration you have set up. This includes various RDS roles such as RD Licensing, RD Session Host, RD Connection Broker, and RD Gateway as well as the vSphere HA and FT settings. Test your proposed availability scenario by forcing a failure and ensuring that you meet or exceed your target RPO/RTO objectives and that your user data is safe and recoverable. If you plan to run these RDS virtual machines alongside other virtual machines in a vSphere cluster, ensure that the proposed consolidated environment will meet or exceed production requirements established by your SLAs if appropriate.

Monitoring the Virtual RDS Farm

Managing the virtual RDS runtime environment is not different from monitoring the physical environment. RDS administrators typically use the Microsoft Windows performance monitor tool Perfmon to collect statistics. RDS integrates with Perfmon to provide familiar counters that indicate system performance. RD Session Host is the more resource-intensive role, so you should be monitoring it closely. You can continue to follow the usual RDS-specific performance counters as you would in a physical environment. Refer to Microsoft Remote Desktop Services best practices for details.

You should use the VMware-provided tools to monitor the vSphere environment:

- vCenter
- `esxtop/resxtop`

Pay close attention to the following VMware vSphere counters (Table 9.2), referring to online documentation on performance monitoring and analysis for more information on these counters and how to interpret them.

Table 9.2 lists a few key counters you should add to your list of inspection points. Of the CPU counters, the total used time indicates system load. Ready time indicates overloaded CPU resources. A significant swap rate in the memory counters is a clear indication of a shortage of memory, and high device latencies in the storage section point to an overloaded or misconfigured array. Network traffic can also cause RDS performance problems. Check total throughput on the NICs to see whether the network is saturated on the vSphere host.

TABLE 9.2: VMware vSphere Performance Counters of Interest to RDS Administrators

SUBSYSTEM	ESXTOP COUNTER	VCENTER COUNTER
CPU	%RDY	Ready
	%USED	Usage
Memory	%ACTV	Active
	SWW/s	Swapin
	SWR/s	Swapout
Storage	ACTV	Commands
	DAVG/cmd	deviceWriteLatency
	KAVG/cmd	deviceReadLatency
		kernelWriteLatency
		kernelReadLatency
Network	MbRX/s	packetsRx
	MbTX/s	packetsTx
	%DRPTX	droppedTx
	%DRPRX	droppedRx

Monitor I/O latency statistics and any disk I/O queuing. Use `esxtop` (or `resxtop`) to see reports of device latency, time spent in the kernel, and latency seen by the virtual machine. You can also monitor I/O queuing in the VMkernel by checking the number of queued commands.

Summary

In this chapter, we reviewed Remote Desktop Services roles and enhancements in Windows Server 2008 R2. We covered the considerations for virtualizing these RDS roles. Except for the RD Session Host role, the RDS roles are usually lightweight services and can be easily virtualized. The RD Session Host server is a multiuser environment whose performance depends upon the number of users and application instances, among other things. It is a challenging server application to virtualize; however, using recently available processors that support hardware MMU virtualization can significantly improve your virtualization experience.

We also discussed the steps to characterize your RDS deployment, check for any application compatibility issues, and define user requirements in an RDS farm. We provided guidelines to size your RDS environments and discussed several deployment considerations when implementing the various RDS roles.

Finally, we followed the fictitious company Source Discovery along its journey of a total Microsoft data-center virtualization. We deployed a proof-of-concept RDS farm consisting of several virtual machines. We discussed virtual machine configuration for each server role in the

proposed farm. Source Discovery published Microsoft Office–configured RemoteApps to publish applications. We also implemented a load-balanced RD Session Host Server farm using RD Connection Broker Load Balancing and high availability for the RD Gateway using Windows Network Load Balancing. We also leveraged vSphere DRS and HA to further minimize any planned and unplanned downtimes for the virtualized RDS farm.

After careful evaluation of this virtualized RDS implementation, Source Discovery will transition the proof of concept into production. Source Discovery also plans to gradually expand the applications running on RD Session Host beyond Microsoft Office and also include more line-of-business applications used by their employees.

Index

Note to the Reader: Throughout this index **boldfaced** page numbers indicate primary discussions of a topic. *Italicized* page numbers indicate illustrations.

Numbers

32-bit vs. 64-bit
  Active Directory domain controllers, **141–142**
  RD Session Host role and, 421–422
  SQL Server VM, **317–318**
  Windows architecture, 418

A

access controls, securing domain controller, **149–150**
access points, failover clusters, *133*, **133**
ACID properties, SQL Server VMs, 314
Active Directory
  32-bit or 64-bit AD domain controllers, **141–143**
  back up procedures for domain controllers, 62
  client certificates for IIS 7, **199**
  configuring RDS VMs, **434**
  HUB servers and, 260
  HUB servers using for messaging, 241
  preparing, *282*
  preparing directory environment, **279–281**, *280*
  preparing domains, **282–284**, *283*
  preparing schema, *281*, **281–282**
  readying guest OS for deployment, **113**
Active Directory Administrative Center (ADAC), 279
Active Directory Application Mode (ADAM), **31**
Active Directory Best Practices Analyzer. *See* ADBPA (Active Directory Best Practices Analyzer)
Active Directory Users and Computers (ADUAC) container, 171
Active Directory VMs
  backup and recovery, **147–149**, *148*
  characterizing workload, **140–141**
  host placement and configuration, **143**
  monitor replication, **150–154**, *151–154*
  overview of, 139
  secure domain controller VM files, **149–150**
  summary review, 176–177
  time synchronization, **144–147**, *145*, *147*
Active Directory VMs, proof of concept
backup Active Directory domain controllers, **168–171**, *169–171*
  configuring network, *159*, **159**
  configuring/synchronizing time, **165–168**, *165–168*
  installing Active Directory service, *160–164*, **160–165**
  installing Tools, *158*, **158–159**
  installing Windows Server 2008 R2 for, **156–158**, *156–158*
  monitoring replication, **173–176**, *173–176*
  overview of, *155–156*, **155–156**
  preventing deletions, **171–172**, *172*
Active Manager, Exchange Server, 249
ADAC (Active Directory Administrative Center), 279
ADAM (Active Directory Application Mode), **31**
ADBPA (Active Directory Best Practices Analyzer)
  domain controller backups, **169–171**, *169–171*
  installing Active Directory service, **163–164**, *164*
  monitoring Active Directory replication, 151, **174–176**, *175–176*
  preventing Active Directory deletions, *171–172*, **171–172**
  scanning roles, **165**
Add Features Wizard, ESX for Windows NLB, **215–216**, *215–216*
Add Hardware Wizard, ESX for Windows NLB, **212–213**, *212–213*
Add Host to Cluster, ESX for Windows NLB, *221*, **221**
Add Roles Wizard
  installing Active Directory service, **160**, *160*
  installing IIS, *202–205*, **202–205**
  installing IIS Remote Management, **207–209**, *207–209*
administration. *See* management services
ADUAC (Active Directory Users and Computers) container, 171
Advanced Options page
  building failover cluster, **124**, *124*
  creating new VM, *98*, **98**
Agent, vCenter Converter, 117
alerts, for thin provision systems, **97**
AMD. *See* processors
anonymous authentication, IIS 7, **199**
anti-affinity rules, DRS, *351–352*, **351–352**
antivirus tools. *See* virus controls
APIs (application programming interfaces)
  conforming backup products in virtual environment to, *2*
  VMSafe, **24–25**
  vStorage, **11**, **14**
`appcmd.exe` tool, IIS installation, **182**
application inventory, RDS deployment and, 426
application programming interfaces. *See* APIs (application programming interfaces)
application server role, SharePoint Server, *375*, **377**
application services
  availability, **19–23**
  fault tolerance, **23**, *23*
  managing/administering with vCenter Server, **28**
  overview of, *2*, 19
  scalability, **25–27**
  security, **24–25**
application-aware availability, **50**
applications
  capacity planning for SharePoint servers, *383*
  checking application compatibility, 427, **427–428**
  diversity in RD Session Host server role, **421–422**
AppSpeed Analysis and Topology Map screen, **232**, *233*

AppSpeed Server
configuring, *229–234*, **229–234**
deploying, **224–229**, *224–229*
installing, *189*, **189**
monitoring performance in IIS VMs, **188–189**, *189*
overview of, **35–36**
performance monitoring in SQL Server VMs, *369*, 369
authentication
IIS 7, **199**
RD Gateway server role, *468–469*
RDS Session Host servers, *443*, 443
authorization, RD Gateway server role, *468*, *470*
Auto Attendant feature, UM server role, 244
AutoIT, 484
automation levels, DRS, **9**, *10*
availability
CAS arrays providing, 240
data redundancy and, **248–249**
hardware, **247**
planning virtual deployment, 45
requirements for SQL Server VMs, **331**
service, **247–248**
availability services
Data Recovery, **23**
high-availability, *22*, **22–23**
for planned hardware downtime, **19–20**
Storage VMotion, **21**, *22*
for unplanned hardware downtime, **20**
VMotion, **20–21**, *21*

B

back pressure, performance monitoring in Exchange Servers, 253
back-end server, Exchange Server 2003, 237
backup and restore. *See also* disaster recovery
in Active Directory VMs, **147–149**, *148*
data protection and, **134–136**, *135–136*
in Exchange Server VMs, **274–275**
Site Recovery Manager automating, **36–38**
in SQL Server VMs, **325–326**, **331–332**, 335–336
third-party solutions for, **136**
using Data Recovery, **20**, **23**, **135–136**, *136*
using Windows Server Backup, **134–135**, *135*
backup destination
domain controllers, **170**
formatting, **61–62**, *61–62*
performing backup job, 62
VDR appliance, **58–59**, *59*
backup strategy, devising
for AD domain controllers, **168–171**, *169–171*
configuring Data Recovery for, **55–62**, *55–62*
installing Data Recovery for, **52–54**, *53–54*
overview of, 51
for tier-1 application, **62–65**, *63–65*
using Data Recovery for, **52**
bandwidth, monitoring IIS VMs, **186**
bandwidth throttling, for IIS websites, **192**, *192–193*
bare metal (or type 1 native) virtualization, **3**, **43**
baselines/baseline groups
attaching to vSphere objects, *74*, **74**
creating for patches with Update Manager, **68–73**, *69–73*
remediating selected vSphere objects, **75–78**, *76–78*
scanning selected vSphere objects, **74–75**, *75*
basic authentication, IIS 7, **199**
benchmarks, Transaction Processing Council, 362
bidirectional traffic shaping, vDS enabling, **18**
binary translation technique, virtualization, **3**, **44**
Boot CD, vCenter Converter, 117
boot controller, PVSCSI controller as, **95**
boot device, manually choosing, *102*, **102**
bridgehead server, Exchange Server 2003, 237
buffer pool pages, locking, *346*, **346–347**
business asset inventory, RDS deployment, 426
businesses, vSphere editions for, **38**, **39**

C

cache
locking SQL Server buffer pool pages in memory, *346*, **346–347**
monitoring memory in IIS, **185**
optimizing IIS websites with output, **193**, *193*
CALs (client access licenses)
overview of, **47–48**
RD licensing role, 422–423
RDS VMs and, 437–439
CAP (Connection Authorization Policy), RD Gateway, *469*
capacity planning
assessing IT infrastructure, *46*, **46–47**
capturing resource utilization baselines, 328, *329*
sizing RDS deployment, 430–431
SQL Server VM, **332**
CapacityIQ, vCenter, **35–36**
CAS (Client Access server) role
array, *240*
building virtual machine for, **284–289**
combining with Hub Transport server, **261–262**
deploying, **290–293**, *291–292*
functions of, 255–256
memory requirements, 257
overview of, **238–241**
processor configurations, 256
storage in, 257–258
CAS/Hub Transport combined server role, 261–262
CCR (Cluster Continuous Replication), Exchange 2007, **23**, 248–249
Central Site Administration VM, configuring, **402–404**, *403–404*
Chargeback, vCenter, **36–37**
CIB (Cluster in a box)
options for building Microsoft clusters, 396
vSphere clustering configurations, 324
CIM (Common Information Model), hardware monitoring with, *32*, **32**
Cisco Nexus 1000V Series virtual switch, vDS support for, **18**
CLI (command-line interfaces)
configuring domain controller backups, **169–171**, *169–171*
defined, **28**
installing IIS with `appcmd.exe` tool, **182**, *182*
overview of, **35**

Client, vCenter Converter, 117
client access licenses. *See* CALs (client access licenses)
Client Access server role. *See* CAS (Client Access server) role
Client Windows Guest Customization Wizard, **110–114**, *111–114*
Clone To Template, VM context menu under Template, **108–110**
cloning, provisioning RDS VMs, 432
Cluster across boxes (CAB), 325, 396
cluster compliance checks, vCenter built-in, **31**
Cluster Continuous Replication (CCR), Exchange 2007, **23**, 248–249
Cluster Installation Rules, *357*, **357**
Cluster IP Addresses dialog box, ESX for Windows NLB, *219*, **219**
Cluster Operations mode, Cluster Parameters, **220**, *220*
Cluster Parameters dialog box, ESX for Windows NLB, **220**, *220*
clusters. *See also* failover clustering
  choosing for VDR, **53–54**, *54*
  defined, **118**
  Exchange Server solutions, 248–249
  high-availability, **23**
  load balancing. *See* WNLB (Windows Network Load Balancing)
  MSCS. *See* MSCS (Microsoft Cluster Service)
  VM placement in vSphere, **81–82**
  WFC (Windows Failover Clustering). *See* WFC (Windows Failover Clustering)
command-line interfaces. *See* CLI (command-line interfaces)
communication protocols, CAS server supporting, 255–256
community training area, 374
compatibility
  building failover clusters, **123**, *124*
  issues in RDS, 427
  vSphere guides for, **39–40**
  Windows 2008 editions as guest OS, **86**
Compliance view, Update Manager, *75*, **75**
compression, HTTP, **193–195**, *194*
concurrent users
  application compatibility in RDS and, 427
  impacting mailbox performance, 268
Configuration Complete screen, VDR, **59**, *59*
Confirm Installation Selection screen, Add Roles Wizard for IIS, *205*, **205**
Connect to a Server, IIS Manager, **209–211**, *210*
connection, to VDR application and appliance, **56–57**, *57*, *61*, **61**
Connection Authorization Policy (CAP), RD Gateway, 424, *469*
consoles
  installing Windows 2008 R2, *101*, **101**
  Tools for, **158**
  VDR, *55*, **55–56**
consolidation process
  Consolidated Backup, 325
  Guided Consolidation, **27**, **33–34**
  SQL Server VM options, *315*, **315–316**
content conversion, CAS servers, 257–258
content expiration, optimizing IIS websites, *191*, **191**
content types area, SharePoint governance plan, 374
continuous optimization, DRS resource allocation, **9**
convergence, (Windows Network Load Balancing), 412
Convert To Template, VM context menu under Template, **108–109**
Converter
  creating virtual machines using, **48**, **116–117**
  migrating to RDS VMs, 431–432
  migrating to SQL Server SMs, 335–336, *336*
converting physical to virtual machine, Windows 2008 VMs, **116–118**
counters
  capturing resource utilization baselines, 329–331
  performance monitoring for UM servers, 272–273
CPU scheduler, ESX, **5**
CPUs. *See* processors
Create Cluster Wizard, *132–134*, **132–134**, *353*
Create New Virtual Machine Wizard, *91–93*, **91–93**
credentials
  configuring VDR, **57–58**, *58*
  connecting to remote server, **210**
  installing Active Directory service, **162**, *162*
Ctrl+Alt+Del screen, and Sysprep, **116**
Customization Specifications Manager
  deploying VM from template, **115**, *115*
  readying guest OS for deployment, **110**, *111*, **114**
Customize Using An Existing Customization Specification, deploying VM from template, *115*, **115**
Customize Using the Customization Wizard, deploying VM from template, **114–115**, *115*

D

DAGs (Database Availability Groups)
  application-aware availability using, **50**
  creating, **304–306**, *305–306*
  high-availability using, **49**
  Mailbox servers deployed as, 243–244
DAS (direct-attached storage), Exchange Server 2010, 277
data integrity, SQL Server VMs and, **314**
data privacy, issues in RDS, 427
data protection
  Data Recovery, **135–136**, *136*
  vStorage APIs for, **14**
  Windows Server Backup, **134–135**, *135*
Data Recovery. *See* VDR (Data Recovery)
data redundancy, **248–249**
data replication, impacting mailbox performance, 268
Database Availability Groups. *See* DAGs (Database Availability Groups)
database back-end server role, SharePoint Server
  configuring, *397–402*
  responsibilities of, 375
  workload related to, **378**

database back-end VM, building/configuring, **396–402**
database cache, Mailbox servers and, 263, 266–267
database consolidation options, SQL Server VMs, *315*, **315–316**
database files
  placement for Active Directory VMs, **163**
  separating Active Directory log files from, **157**
database mirroring
  application-aware availability using, **50**
  migrating SQL Servers to virtual environment, 337
databases, Exchange Server 2010
  copying, **308–309**, *309*
  creating, **306–308**, *307–308*
  Mailbox servers and, 243–244
Datacenter version of Windows Server 2008, choosing for IIS VM, 180
datastore
  choosing for VM, **92**, *92*
  Deploy OVF Template Wizard, *227*, **227**
  storing templates in, **108**
DCDiag (Domain Controller Diagnostics) tool
  monitoring Active Directory replication, 151
  overview of, *152–153*, **152–153**
  for Windows Server 2003, 154
decision support system (DSS), database workloads, 327
deduplication stores, Data Recovery, **23**, **135–136**
deletions, preventing in Active Directory VMs, **171–172**, *172*
Deploy From File, Deploy OVF Template Wizard, *224*, **224**
Deploy From Template Wizard, Exchange Server 2010, *285–286*
Deploy OVF Template page, VDR installation, 52
Deploy OVF Template Wizard, AppSpeed installation, *224–228*, **224–228**
Deploy Virtual Machine From This Template, **110**, **114**
deployment
  Internet Information Services, **180–182**, *182*
  SQL Server VMs, **333–335**, **360–362**, *361–362*
  Windows failover cluster. *See* WFC (Windows Failover Clustering), deploying
deployment, Exchange Server VMs
  CAS server role, **290–293**, *291–292*
  Hub Transport server role, **293–296**, *294–295*
  Mailbox server role, **296–298**, *297–298*
deployment, Microsoft Server applications
  patching with Update Manager, *65*, **65–66**
  planning virtual, **44–47**, *46*
  using scale-up or scale-out strategy, **81**
  VM placement in vSphere cluster, **81–82**
deployment, RDS VMs
  considerations, **431**
  DRS and HA considerations, **432–433**
  initial creation of RDS VMs, **431–432**
  inventories required, **425–427**
  monitoring farm, **484–485**
  ongoing provisioning strategy, **432**
  single-server vs. farm models, **425**
  sizing, **429–431**
  testing farm, **480–484**
deployment, Windows 2008 VMs
  choosing virtual hardware, **88–90**
  converting physical machine to virtual machine, **116–118**
  creating virtual machine, *90–100*, **90–100**
  installing Windows 2008 R2, *101–106*, **101–106**
  overview of, **87–88**
  synchronizing time, **107–108**
  using templates, *108–109*, **108–116**, *111–116*
deployment, Windows failover clusters
  building failover cluster, **120–127**, *122–127*
  clustering VMs, **118–119**
  configuring, **127–134**, *127–134*
  hardware/software requirements, **119–120**
  overview of, **118**
detach and attach approach, migrating SQL Servers to VMs, 337
DHCP (Dynamic Host Configuration Protocol), *159*
digest authentication, IIS 7, **199**
disaster recovery
  for Active Directory VMs, **149**
  planning virtual deployment, 45
  using Data Recovery, **20**, **23**
  using database availability groups, **49**
  using Site Recovery Manager, **36–38**
Disk Management, 130
Disk Manager
  failover cluster configuration, *128*, **128**
  formatting partitions for SQL Server VMs, 344, *344*
disks (or partitions)
  alignment, **129**
  alignment inside a Windows VM, 343
  clustered SQL Server VMs, 348
  compatibility mode, **124**
  configuring for SQL Server VMs, **348–350**, *349–350*
  converting physical machine to VM, **117**
  creating thin-provisioned virtual, **12–13**, *13*
  creating/configuring Active Directory VM, 155, **157**
  failover cluster configuration, **128–131**, *128–131*, **356**, *356*
  monitoring performance in IIS VMs, **185–186**
  prerequisites for SharePoint, 377
  resource management, **6–7**
  selecting for new virtual machine, **96–98**, *97*
  SQL Server VMs, 340
  VMFS vs. RDM for SQL Server VMs, 322–323, *323*
Distributed Management Task Force (DMTF), SMASH profiles defined by, **32**
Distributed Power Management. *See* DPM (Distributed Power Management)
Distributed Resource Scheduler. *See* DRS (Distributed Resource Scheduler)

Distributed Switch, VMware vNetwork, **15**
DMTF (Distributed Management Task Force), SMASH profiles defined by, **32**
DMZs (demilitarized zones)
EDGE server and, 238
firewalls and, **199–200**, *200*
vShield zones and, **200**
DNS (Domain Name Service), Active Directory workload, 140
DNS round-robin, IIS server, *195*, **195–196**
Do Not Customize, deploying VM from template, **114**, *115*
domain controllers, Active Directory VMs
64-bit vs. 32 bit, **141–142**
backup and recovery, **147–149**, *148*
backup process, **168–171**, *169–171*
characterizing workload, 140
configuring time synchronization, **107**
host placement/ configuration, **143**
installing Active Directory service, *160–164*, **160–165**
monitoring Active Directory replication, **150**
network configuration, **142–143**, *159*, **159**
securing virtual machine files, **149–150**
sizing, *155–156*, **155–156**
time synchronization, **144**, *145*
Domain Name Service (DNS), Active Directory workload, 140
Domain Secure Channel (NlTest), 151, **153–154**, *155*
downtime area, SharePoint governance plan, 373–374
DPM (Distributed Power Management)
built-in cluster compliance check for, **31**
defined, **2**
overview of, **10**, *10*
drivers, installing Windows 2008 R2, *103–104*, **103–105**
DRS (Distributed Resource Scheduler)
administering, **28**
cluster compliance checks for, **31**
creating anti-affinity rules, *351–352*, **351–352**
enabling for vSphere clusters, **81–82**, 338, *338*
high-availability settings, *22*, **22**, 324
improving scalability with, **25–26**
installing AppSpeed with, **226**, *226*
RDS VMs and, **432–433**, 433
for smaller virtual machines, **81**
as vCompute service, **2**, *8–10*, **8–10**
VMFS as key enabler for, **11**
DSS (decision support system), database workloads, 327
dynamic baselines
automatic updates of, **68**, *69*
remediating selected vSphere objects, **76**, *76*
Dynamic Host Configuration Protocol (DHCP), *159*

E

eagerzeroedthick, VM disk files, **97**
Edge Sync process, Hub Transport server, **258**
Edge Transport server role (EDGE)
functions, 250–251
overview of, **237–238**, *239*
processor configurations, 251–252
storage in, 252–255
Edit Output Cache Settings dialog box, *193*, **193**
editing
VM configuration, *60*, **60**
VM settings prior to creation, **99**, *99*
email, UM servers and, 270
EMC (Exchange Management Shell)
copying mailbox databases with, **308–309**
creating DAGs with, **304–306**
creating mailbox databases with, **306–308**
moving Exchange Server 2003 mailboxes to Exchange Server 2010, **309–311**
end user license agreement. *See* EULA (end user license agreement)
Enhanced VMotion Compatibility (EVC), **20**
Enterprise Windows Server 2008, for IIS VM, 180
environment usage area, SharePoint governance plan, 373
EPT (Extended Page Tables)
defined, **5**
hardware-assisted memory management, **44**
Intel processor technology, 420
ESX/ESXi servers
advanced memory management, **6**
best practices for configuring, 338–340
configuring for Windows NLB, *211–222*, **211–222**
configuring hosts for Windows NLB, **222–223**, **222–234**, *223*
data integrity of SQL Server VMs and, 314
defined, **3**
Distributed Power Management, *10*, **10**
Distributed Resource Scheduler, **8–10**, *8–10*
failover clustering unsupported in mixed environments, **120**
hardware-assisted memory virtualization, **4–5**
Host Update Utility, **35**
Management Assistant for, **34**
managing/administering, **28**
as new generation of virtual hardware, *7–8*, **7–8**
patching/upgrading with Update Manager, **34**
processor scheduling, **5**
relocating SQL Server VMs to other, 332
resource management, **6–7**
scalability, **25–26**
shared storage scalability, 322
SQL Server VMs dedicated to, 334
validated in Microsoft SVVP program, **79**
as vCompute service, **2**
virtualization technologies, **3–4**, *4*
vStorage services for. *See* vStorage services
vswif network adapter for, **16**
`esxtop` command tool
monitoring in I/O latencies SQL Server VMs, 366
monitoring RDS virtual farms, 484
monitoring SQL Server VMs, *363*, 363
performance counters, 485
viewing ESX server workload, *187*, **187–188**

Ethernet, creating new VM, **95**, *96*
Ethernet adapter, configuring ESX for Windows NLB, *212–213*, **212–213**
EULA (end user license agreement)
  Deploy OVF Template Wizard, *225*
  installing Exchange Server 2010, 291
  installing SharePoint software, 407
  installing Windows Server 2008 R2 for Active Directory, **157**
EVC (Enhanced VMotion Compatibility), **20**
Exchange Management Shell. *See* EMC (Exchange Management Shell)
Exchange Server 2003
  comparing Exchange versions, 245–246
  moving mailboxes to Exchange Server 2010, **309–311**
  server roles, **237**
Exchange Server 2007, **237–238**, 245–246
Exchange Server 2010 VM
  backing up and restoring data, **274–275**
  Client Access server role, **238–241**, *240*
  comparing Exchange versions, 245–246, **245–247**
  data redundancy and availability, **248–249**
  dual-homed UM virtual machine, *270*
  Edge Transport server role, **237–238**, *239*
  Exchange Server 2003 and, **237**
  hardware availability, **247**
  Hub Transport role, **241–242**, *242*
  Mailbox server role, **242–244**, *243*
  overview of, 235
  planning virtual deployment, **44**
  Requirements Calculator, *269*
  server roles, **237–238**
  service availability, **247–248**
  summary, 311
  trends in, 236–237
  Unified Messaging server role, **244–245**
  VMware building blocks for, **275–276**
  workloads and, **250**
  workloads of Client Access and Hub Transport combined server, **261–262**
  workloads of Client Access server, **255–258**
  workloads of Edge Transport server, **250–255**
  workloads of Hub Transport server, **258–261**
  workloads of Mailbox server, **262–270**
  workloads of Unified Messaging server, **270–274**
Exchange Server 2010 VM, proof of concept
  building CAS role virtual machine, **284–289**
  building environment for, 278–279
  Database Availability Groups and, **304–306**, *305–306*
  database copies, **308–309**, *309*
  database creation, **306–308**, *307–308*
  Deploy From Template Wizard, *285–286*
  deploying CAS role, **290–293**, *291–292*
  deploying Hub Transport role, **293–296**, *294–295*
  deploying Mailbox role, **296–298**, *297–298*
  installation options, *289*
  installing Exchange Server 2010, *291*
  installing prerequisites using scripts, *288*
  licensing servers, **302–304**, *303–304*
  LoadGen 2010 Beta for testing, **298–302**, *299–302*
  moving Exchange Server 2003 mailboxes to, **309–311**, *310–311*
  overview of, 277–278
  preparing Active Directory, *282*, **282–283**
  preparing Active Directory domains, *283*, **283–284**
  preparing Active Directory schema, *281*, **281–282**
  preparing directory environment, **279–281**, *280*
  `Scripts` directory, *287*
  security groups, *284*
  Setup program, *290*
Exchange Server Profile Analyzer, planning virtual deployment, 46
Extended Page Tables. *See* EPT (Extended Page Tables)

F

failover clustering. *See also* MSCS (Microsoft Cluster Service); WFC (Windows Failover Clustering)
  application-aware availability and, **50**
  configuration, **127–134**, *128–134*
  high-availability and, **49**, 323
  installing, **130–131**, *131*
  vSphere configurations for, 324–325
failover clustering, in SQL Server VMs
  Cluster Installation Rules, **357**, *357*
  configuring, 350–351
  creating and validating failover clusters, *353–354*, **353–354**
  disabling VM restart priority, **352**, *352*
  DRS anti-affinity rules, 350–351, *351–352*
  first node of failover cluster successfully installed, **357**, *357*
  installing SQL Server in failover cluster, **355–356**, *355–356*
  Microsoft Distributed Transaction Coordinator and, **354**
  second node of failover cluster successfully installed, *358–359*, **358–359**
  selecting disk resources for failover cluster, **356**, *356*
  testing failover cluster services, **359–360**, *360*
farm mode, distributed model for SharePoint servers, 374
Fault Tolerance. *See* FT (Fault Tolerance)
fax communication, UM servers, 270

Fibre Channel SANs
deploying clustered VMs using, **119**
Storage VMotion working across, **21**
vs. iSCSI for SQL Server, 322
firewalls, IIS 7 security, **199–200**, *200*
FLR (File Level Restore) Client feature, Data Recovery, **135**
format warning, chosen backup destination, *61*, **61**
forms authentication, IIS *7*, **199**
front-end server, Exchange Server 2003, 237
FSMO roles
DCDiag checking Active Directory status, 152
time synchronization using, *168*, **168**
virtualizing Active Directory, 161
FT (Fault Tolerance)
administering with vCenter Server, **28**
built-in cluster compliance checks for, **31**
defined, **20**
failover clustering unsupported in VMs with, **120**
high-availability features in VMware, 324
overview of, *23*, **23**
protecting VMs with, **118**
VM disk file requirements, **97**, *98*
VMFS as key enabler for, **11**
Fusion, 43

G

GC (global catalog) servers
Active Directory workload and, 140
installing Active Directory service, **162**
golden image, **108**
golden template, **88**
GOS (guest OS)
backup solutions, **134–135**, *135*
compatibility and requirements, **86–87**
creating to synchronize time, **107–108**
deploying clustered VMs, **119**
deploying VMs from template, **114–116**, *115–116*
determining virtual hardware settings, **93**, *94*
readying guest OS for deployment, **110–114**, *111–114*
governance, SharePoint. *See* SharePoint governance
GPO (Group Policy Object)
modifying Windows Time Service, *147*, **147**
time synchronization with, **146**, *166*, **166–168**
GPS clock, time synchronization, **144**
guest customization, deploying VMs using templates, **110–114**
Guest Customization page, template deployment, **114**, *115*
Guided Consolidation, vCenter, **27**, **33–34**

H

HA (high-availability). *See also* HA (high-availability)
administering, **28**
cluster compliance checks for, **31**
clustering with Windows Server 2003/2008 and, **23**
defined, **20**
failover clustering unsupported in VMs with, **120**
in Microsoft Server VMs, **49–51**, *50*
overview of, *22*, **22**
provided by clusters, **118**, 338, *338*
in RD Connection Broker servers, 455
in RDS VMs, **432–433**, 433
in SQL Server VMs, **323–325**
VMFS as enabler for, **11**
HA (high-availability), in Exchange Server
data redundancy and availability, **248–249**
hardware availability, **247**
Mailbox servers and, 243
overview of, 247
service availability, **247**
hardware
availability and. *See also* availability services, **247**
capacity planning for SharePoint servers, *383*
for Central Site Administration VM, 403
for database back-end VMs, 397
ESX for Windows NLB, **212–213**, *212–213*, *215*, **215**
for failover clusters, **119**
for IIS VMs, **180**, **182–183**, **196**
for index server VMs, **404–405**
for Microsoft Server VMs, **44**
monitoring with CIM, *32*, **32–33**
new generation of virtual, *7–8*, **7–8**
planning virtual deployment, 45
for RD Connection Broker servers, 453
for RD Gateway server role, 466
for RD Licensing role, 434
for RD Web Access server role, 450
for RDS deployment, 427
for RDS Session Host servers, 440
for SQL Server VMs, **321**, *332*, 332
using firewalls, **199**, *200*
vStorage improvements for iSCSI, **15**
for WFE server VMs, **405–406**
for Windows Server 2008 VMs, **88–90**
hardware acceleration, installing Windows 2008 R2, *106*, **106**
Hardware Status plug-in, **32**
hardware-assisted virtualization
for memory, **4–5**, *421*
overview of, **3–4**, *4*
heartbeat network
cluster node 1 configuration, **121**
cluster node 2 configuration, **126**
cluster node build specification, 120
deploying clustered VMs using, **119**
failover cluster configuration, **133**
host availability, **50**
host profiles, **29–30**, *30–31*
Host Update Utility, **35**
hosted (type 2 or hosted hypervisor) virtualization, **3**, *42*, **42–43**

hosts
configuring Active Directory VMs, **142–143**
configuring ESX hosts for Windows NLB, **222–223**, *223*
creating virtual machines at, 91
implementing Active Directory VMs, **143**
not deploying SQL Server VMs with, 314, *314*
patch management using, **34**
planning virtual deployment, 45
Update Manager deployment options, **65–66**
hot extend for virtual disks, **12**, **26**
hot plug devices, **7**
hot-add support, **7**, **25–26**, *27*
HP tool
Build A Solution option, *387*
Business Profile page, *388*
choosing SharePoint Server architecture, *392–393*
customizing workload mix, *390*
entering company profile information, 386, *386*
selecting usage pattern for SharePoint, *389*
sizing SharePoint servers, *386–395*, **386–395**
Solution Alternatives page, *395*
Solution Reference page, *388*
solution results details, *394*
usage characteristics page, *391*
HTTP compression, **193–195**, *194*
HTTP keep-alive responses, optimizing IIS, **191–192**
HTTP Response Headers, optimizing IIS, *191*, **191**
HTTP Secure (HTTPS), RD Gateway role using, 423
HTTPS (HTTP Secure), RD Gateway role using, 423
`http.sys` HTTP listener, tuning IIS, **190**
Hub Transport server role
combining with CAS server role, **261–262**
deploying, **293–296**, *294–295*
Edge Sync process, **258**
functions of, 258
overview of, **241–242**, *242*
transport dumpster, **258–261**
Hyper-V, RDVH integration with, 416
hypervisor architecture
SQL Server VM deployment, 314, *314*
type 1 native or bare-metal, **43**
type 2 or hosted, *42*, **42–43**

I

icon change, cloning VM to template, **109**
IDE virtual devices, **7**
IIS (Internet Information Services)
as prerequisite for SharePoint, 376
RD Gateway server role and, *471*
RD Web Access server role and, 450–451, *451*
IIS (Internet Information Services) VMs
choosing flavor of Windows Server 2008, **180**
installing IIS, **181–182**, *182*, *202–211*, **202–211**
network load balancing, **195–198**
overview of, 179
performance, **182–189**
performance tuning, **189–195**
securing, 198–201
summary review, 234
IIS (Internet Information Services) VMs, proof of concept
installing AppSpeed, *224–234*, 224–234
installing IIS, 202–211, *202–211*
network load balancing, 211–223, *211–223*
overview of, 201–202
IIS 7.0 Administrator's Pocket Consultant, 182
IIS Manager
configuring bandwidth throttling, **192**, *192–193*
connecting to remote server, **209–211**, *209–211*
installing IIS, **181**
installing IIS Remote Management service, *208*, **208**
launching from Administrative Tools, *206–207*, **206–207**
IMAP (Internet Message Access Protocol), CAS server supporting, 255–256
index server role, SharePoint Server
configuring, **404–405**, *405*
responsibilities of, 375
workload related to, **377**
infrastructure services
managing/administering, **28**
overview of, **1–2**, *2*
planning virtual deployment, 45
vNetwork, **15–19**, *16*
vStorage, **11–15**, *13*
initializing disk, failover cluster configuration, *128*, **128**
Input/Output per second. *See* IOPS (Input/Output per second)
Install Now button, Windows 2008 R2, *101*
Installation Results screen, Add Roles Wizard for IIS, **205**, *205*
Intel. *See* processors
Internet Information Services. *See* IIS (Internet Information Services)
Internet Message Access Protocol (IMAP), CAS server supporting, 255–256
Internet Security and Acceleration (ISA) server, **197**
inventory listing, backup of tier-1, 62–65
Inventory option, configuring VDR, 56
inventory search
vCenter Server 4, **28–29**, *29*
vCenter Server Linked Mode, **31–32**
I/O devices
monitoring in IIS VMs, **185–186**
monitoring latencies in SQL Server VMs, 366
monitoring queuing in SQL Server VMs, 366
patching with Update Manager, **66**
VMDirect Path I/O for Networking accessing, **19**
I/O enhanced performance
PVSCI storage adapter for, **13–14**
VMDirectPath I/O for storage with, **15**
VMXNET3 network adapter for, **15–16**, *16*
vStorage improvements, **14**
IOMeter tool, 361–362, *362*

IOPS (Input/Output per second)
message rates in EDGE servers, 255
message rates in HUB servers, 259
message rates in MX servers, 263, 266
IP addresses
configuring Active Directory VMs, *159*, **159**
configuring ESX for Windows NLB, *218–220*, **218–220**
configuring failover clusters, *133*, **133**
network load balancing IIS server, *195*, **195–196**
Windows Network Load Balancing, **197–198**
IPv6, support for network storage, **16**
ISA (Internet Security and Acceleration) server, **197**
iSCSI storage protocol
failover clustering unsupported on, **120**
Storage VMotion working across, **21**
vs. Fibre Channel for SQL Server, 322
vStorage improvements, **14–15**
ISO image
attaching prior to VM creation, **99**, *100*
creating VM, **90**
defined, **90**
IT infrastructure, Capacity Planner assessing, *46*, **46–47**

J

JBOD (just a bunch of disks), Exchange Server 2010 supporting, 277
jumbo frames, Ethernet storage benefits, **14**

K

KCC (knowledge checker consistency) status, Active Directory, 152
keep-alive responses, HTTP, (optimizing IIS), **191**, *191–192*
Kerberos v5, Active Directory time synchronization, **144**
keyboard functions, VMware Tools for, **158**

L

Lab Manager, vCenter, **36–37**
large memory pages, **5**, **6**
licensing
centralized, **33**
Data Recovery, **136**
Exchange servers, **302–304**, *303–304*
Microsoft Server VMs, **47–48**
RD Licensing role, 422–423
readying guest OS for deployment, **112**
SharePoint software, 407
SQL Server VMs, **317**
Windows 2008 R2, **103**
Windows Server 2008 R2 for Active Directory, **157**
Lifecycle Manager, vCenter, **36–37**
Linked Mode, vCenter Servers, **31–32**
Linux kernels, virtualization, **4**
load balancing. *See also* WNLB (Windows Network Load Balancing)
adding domain controller to existing domain, **161**, *161*
CAS arrays providing, 240
Distributed Resource Scheduler and, 20, **26**, 433
IIS VMs, *195*, **195–198**
for RD Session Host servers. *See* RD Connection Broker server role
Load Driver link, installing Windows 2008 R2, *103*
LoadGen 2010 Beta, testing Exchange Server, **298–302**, *299–302*
LoadRunner utility, HP, 484
location
choosing virtual machine, **91–92**, *92*
choosing Windows 2008 R2, *103*, **103**
VM disk file, **97–98**, *98*
Lock Pages in Memory, SQL Server buffer pool, *346*, **346–347**
log shipping
application-aware availability using, **50**
migrating SQL Servers to virtual environment, 337
login
creating virtual machine, 90–91
Data Recovery, **55**
Domain Controller screen for, **164**, *164*
installing Windows Server 2008 R2 for Active Directory, **157**
Management Assistant noninteractive, **34**
readying guest OS for deployment, **112**
logs/logging
Active Directory VMs, **157**, **163**
CAS servers, 257
EDGE server, 252–253
LSI Logic, default primary adapter, 14
LUNs
building failover clusters, **122**, *123*
configuration for SQL Server VMs, 342, *342*
mapping to VMFS volume, **11–12**

M

MA (Management Assistant), vSphere, **34**
MAC addresses, Windows Network Load Balancing, **197**
Mailbox server role
backing up and restoring data in, **274–275**
deploying, **296–298**, *297–298*
designing storage for, 267–269
functions of, 262
memory requirements, 265–267
overview of, **242–244**, *243*
processor configurations, 262–264
storage in, 267–270
mailboxes, migrating Exchange Server 2003 to 2010, **309–311**
maintenance, planned server, **9**
Majority Node Set clusters, **23**
majority node set (MNS) clusters, Exchange Server, 248
management services
management solutions, **35–38**
overview of, *2*, 27
vApps, **35**
vCenter Guided Consolidation, **33–34**
vCenter Orchestrator, **33**
vCenter Server, **28–33**
vCenter Update Manager, **34**
vSphere command-line interfaces, **35**

vSphere compatibility guides, **39–40**
vSphere editions, **38–39**
vSphere Host Update Utility, **35**
vSphere Management Assistant, **34**
management solutions, **35–38**
Management Tools, Exchange Server 2010, 291–292
Maps view, Active Directory VMs, **142**
memory
64-bit vs. 32 bit AD domain controllers and, **141–142**
for Active Directory VM, 155
advanced, **6**
ballooning, **6**
for CAS and Hub Transport combined role, 261–262
for CAS servers, 257
for EDGE servers, 252
for ESX and ESXi scalability, **26**
hardware-assisted, **4–5**
hot-add support for, **26–27**
for IIS VMs, **180**, **184–185**
for Mailbox server, 265–267
for new virtual machine, *95*, **95**
over-commitment, **6**
patching with Update Manager and, **66**
performance monitoring for UM servers, 272
physical to virtual migration and, **48**
for RDS Session Host servers, 440
resource management, **6–7**
for resource pools, **8**
using large memory pages for SQL Server VMs, **347**
VMkernel protection and, **25**
memory management virtualization. *See* MMU (memory management virtualization)
message rates/size, EDGE server, 251
Microsoft Application Verifier, 428
Microsoft Cluster Service. *See* MSCS (Microsoft Cluster Service)
Microsoft Distributed Transaction Coordinator (MSDTC), *354*, **354**
Microsoft FC (Failover Clustering). *See* WFC (Windows Failover Clustering)
Microsoft Office
adding applications to RemoteApps list, 462, *462*
SharePoint Server 2007. *See* SharePoint Server 2007 VM
Microsoft Premier-level support agreement, **81**
Microsoft Regional Licensing Specialist, 47
Microsoft Server application VMs
backup strategy. *See* backup strategy, devising
hardware, **44**
high-availability, **49–51**, *50*
licensing, **47–48**
Microsoft Server Virtualization Validation Program, 78–81, **78–82**
migration to virtual environment, **48–49**
new era of virtualization and, **41–42**
overview of, 41
patching using Update Manager. *See* Update Manager
planning deployment, **44–47**, *46*
summary review, 82–83
technologies, *42–43*, **42–43**
Microsoft Service Center Operations Manager, 272, 274
Microsoft System Center Capacity Planner. *See* System Center Capacity Planner
Microsoft Windows. *See* Windows
migration. *See also* P2V (physical to virtual) migration
of clustered VMs, **120**
enabling storage, **21–22**, *22*
to virtual environment, **48–49**, **335–337**
VMotion enabling live, **20–21**
MIME (Multipurpose Internet Mail Extensions), in CAS servers, 257–258
MMU (memory management virtualization)
AMD and Intel processors, **44**, 317, 321
hardware virtualization, *421*
RD Session Host server role, 420, *421*
MNS (majority node set) clusters, Exchange Server, 248
modular design, IIS 7 security, **198**
monitoring. *See also* performance monitoring
Active Directory VMs, 140
replication traffic, Active Directory VMs, **150–154**, *151–154*, **173–176**, *173–176*
status of backup job, **64**, *64*
thin provision production systems, **97**
mouse functions, VMware Tools for, **158**
MSCS (Microsoft Cluster Service)
high-availability and, 323
support for, **23**
for Windows 2000/2003, **118**
MSDTC (Microsoft Distributed Transaction Coordinator) and, *354*, **354**
multicast mode, Windows Network Load Balancing, **197**
multipathing, vStorage APIs for, **14**
Multipurpose Internet Mail Extensions (MIME), in CAS servers, 257–258

N

native virtualization, **3**, **43**
navigation, vCenter Server 4 simplifying, **28–29**, *29*
.NET Framework, as prerequisite for SharePoint, *377*
Network and Internet Control panel, ESX for Windows NLB, 214, **214**
Network and Sharing Center, ESX for Windows NLB, **214**, 214
network interface controllers (NICs), **142**, 211
Network Load Balancing. *See* NLB (Network Load Balancing)
Network Load Balancing Manager, **217–218**, *217–218*
Network Mapping screen, Deploy OVF Template Wizard, **227**, *227*
Network page, guest OS deployment, **113**
Network Policy Server, RD Gateway server role, *470*
network share, as domain controller backup destination, **170**
Network VMotion, vDS enabling, **17**
networks
capacity planning for SharePoint servers, *382*
configuring Active Directory VMs, **159**, *159*

connection for new VM, **95**, *96*
managing bandwidth, **6–7**
performance monitoring for UM server, 273
RDS Session Host servers, 440
New Customization Specification page, guest OS deployment, **111**, *111*
New Simple Volume Wizard, failover clusters, **129**, *129*
NFS storage protocol
failover clustering unsupported on, **120**
Storage VMotion working across, **21**
vStorage improvements, **14**
NICs (network interface controllers), **142**, 211
NLB (Network Load Balancing)
IIS VMs, **195–198**, *211–223*, **211–223**
software load balancing with, **196**
Windows. *See* WNLB (Windows Network Load Balancing)
NlTest (Domain Secure Channel), 151, **153–154**, *155*
noninteractive login, Management Assistant, **34**
Notify Switches setting, ESX hosts for Windows NLB, **223**
NPIV (N-Port ID Virtualization), **120**
NTBackup, Active Directory VMs, **147–149**, *148*
`NTDS.DIT` file, configuring Active Directory VMs, **143**
n-tier applications, vApp, **35**
NTP servers, and time synchronization, **107–108**

O

objects, vSphere
attaching baselines/baseline groups to, *74*, **74**
remediating selected, **75–78**, *76–78*
scanning selected, *74–75*, **74–75**
offline virtual machines, patching, **65**
OLTP (online transaction processing), 327
OMA (Outlook Mobile Access), CAS server and, 239
on-disk file locking, VMFS, 11
ongoing provisioning strategy, RDS VMs, **432**
online references
AppSpeed, 189
Capacity Planner, 47
Cisco Nexus 1000V Series virtual switch, 18
Failover Clustering and Microsoft Cluster Service setup for VMware, 23
IIS 7.0 Administrator's Pocket Consultant, 182
licensing considerations, 48
management solutions for virtualization, 36–38
Microsoft Solution Accelerators, 46
Microsoft SVVP program, 79
RFC3484, 196
SPEC website, 183
vCenter AppSpeed, 36
vCenter CapacityIQ, 36
vCenter Chargeback, 37
vCenter Converter, 118
vCenter Update Manager Sizing Estimator, 66
virtual switches, 17
virtualizing business applications, 41
vShield zones, 201
vSphere Best Practices, 143
vSphere business edition comparisons, 38
vSphere centralized licensing, 33
vSphere Purchase Advisor, 39
warning on virtualizing Active Directory domain controllers, **175**, *176*
Windows Server 2008 hardware requirements, 182
online transaction processing (OLTP), 327
Open Virtualization Format. *See* OVF (Open Virtualization Format)
operational requirements, planning virtual deployment, **45–46**
Orchestrator, vCenter, **27**, **33**
Outlook Anywhere, CAS server and, 239, 255–256
Outlook Mobile Access (OMA), CAS server and, 239
Outlook modes, impacting mailbox performance, 268
Outlook Web Access, CAS server and, 239, 255–256
output caching, optimizing IIS, *193*, **193**
over-commitment of storage capacity, virtual disk thin provisioning, **12–13**, *13*
OVF (Open Virtualization Format)
installing AppSpeed, **35**, 189, *224–228*, **224–228**
installing Data Recovery, **135–136**, *136*
installing VDR, **52–53**, *53*
OVF Template Details screen, Deploy OVF Template Wizard, *225*, **225**

P

P2V (physical to virtual) migration
building Microsoft clusters, 396
clustering with virtual machines, 325
overview of, **116–117**
using Converter, **48**, 431–432
page tables, for hardware-assisted memory virtualization, **4–5**
PAL (Performance Analysis of Logs), 274
PAM (Primary Active Manager), Exchange Server, 249
paravirtualized Linux guest operating systems, **4**
paravirtualized SCSI. *See* PVSCSI (paravirtualized SCSI) storage adapter
partitions. *See* disks (or partitions)
passwords
installing Windows Server 2008 R2 for Active Directory, **157**
strong, *158*
patch baselines, *72*, **72**, *74*
patches
deploying with Update Manager. *See* Update Manager
SQL Server VM, **332**
PDC (primary domain controller) emulator role
Active Directory time synchronization, **144**, *145*
proof of concept, *166–168*, **166–168**
time synchronization using Group Policy and WMI filters, **146**
time synchronization using Windows Time Service, **146**

Perfmon
capturing resource utilization baselines, 328
monitoring AppSpeed in, **188–189**, *189*
monitoring disk input/output in, **185–186**
monitoring memory utilization in, **184–185**
monitoring network bandwidth in, **186**
monitoring processor utilization in, **184**
monitoring SQL Server VMs in, 362
monitoring VM statistics in, **186–188**, *187*
overview of, **182–183**
pinpointing IIS bottlenecks with, **183**
planning virtual deployment with, 46
performance
application compatibility issues in RDS, 427
hardware-assisted memory virtualization enhancing, **5**
IIS 7 features, **181**
impacts of physical to virtual migration, **48**
Mailbox servers and, 267–268
planning virtual deployment, 45
of SQL Server in virtual environment, **316–317**, *317*
troubleshooting with AppSpeed, **35–36**
vNetwork enhancing, **18–19**
vStorage enhancing, **14–15**
performance, of IIS VMs
monitoring AppSpeed, **188–189**, *189*
monitoring disk input/output, **185–186**
monitoring memory utilization, **184–185**
monitoring network bandwidth, **186**
monitoring processor utilization, **184**
overview of, **182–183**
tuning, **189–195**, *191–194*
VM statistics in Perfmon, **186–188**, *187*
Performance Analysis of Logs (PAL), 274
performance counters
capturing resource utilization baselines, 329–331
monitoring UM servers, 272–273
vSphere, 365, 485
performance monitoring
Exchange Servers, 253
metrics for EDGE servers, 251
RDS farms, **484–485**
SQL Servers, **362–370**, *363–367*, *369*
UM servers, 272–274
per-processor licensing, virtualizing Microsoft applications, **47**
physical compatibility mode, building failover cluster, **123**, *124*
physical to virtual migration. *See* P2V (physical to virtual) migration
pilot phase, virtual deployment, 46
placement, DRS recommendations for VM, **9**, *10*
plug-ins
configuring AppSpeed Server, *230*, **230**
installing VDR, **55–56**
policies
backup retention, 62
DRS resource allocation, **8–9**, *9*
host profile, **29–30**, *30*
licensing, virtualizing Microsoft applications, **47**
security, **24**
traffic-shaping, **18**
POP (Post Office Protocol), and CAS server, 255
Port Rules dialog box, ESX for Windows NLB, *220*, **220**
Post Office Protocol (POP), and CAS server, 255
power management, DPM, *10*, **10**
Power On This Virtual Machine After Creation box, **115**, *116*
Power On, VDR appliance, **55**
PowerCLI, **35**
Primary Active Manager (PAM), Exchange Server, 249
primary domain controller. *See* PDC (primary domain controller) emulator role
privileges, remote offices/branch offices, **143**
probes, AppSpeed, *231–232*, **231–232**
Probes Management screen, AppSpeed, **232**, *233*
processors
building blocks for, 276
configurations for CAS servers, 256, 261
configurations for EDGE servers, 251–252
configurations for Mailbox servers, 262–264
configurations for UM servers, 271–272
creating resource pools, **8**
creating/configuring Active Directory VM, 155
E1000 network adapter, **16**
ESX/ESXi scalability and, **26**
hardware-assisted memory virtualization using, *4*
hardware-assisted virtualization, **4–5**
hot-add support for, **26–27**
Lance PCNet 32 network adapter, **16**
monitoring performance in IIS, **184**
per-processor licensing, **47**
physical to virtual migration and, **48**
RDS Session Host servers, **420–421**, 439
resource management, **6–7**
scale-up or scale-out deployment strategy for, **81**
scheduling, **5**
selecting for new VM, *94*, **94**
SQL Server VMs, 321, 332
virtualization technology, **44**
product key
installing SharePoint software, 407
licensing Exchange servers, 304
Profile Compliance tab, **29–30**, *30*
properties, configuring VDR, **60**, *60*
proportional resource allocation, **6–7**
proportional-share scheduling algorithm, CPU scheduler, **5**
protocol logging, in CAS servers, 257
provisioning
automating with vCenter Lifecycle Manager, **36–37**
developing ongoing strategy for, **49**
selecting for new virtual machine, **97**, *98*

PVSCSI (paravirtualized SCSI)
storage adapter
configuring SQL Server VMs, *341*, 341–342
creating new VMs, **95**
defined, **11**
overview of, *13*, **13–14**
virtual hardware version 7 using, 93

Q

query server role, SharePoint Server, 375

R

RAID, SharePoint servers and, 378
RAM. *See also* memory, 276
RAP (Resource Authorization Policy), RD Gateway and, 424
Rapid Virtualization Indexing (RVI), AMD, **5**, **44**, 420
RARP packet transmission, ESX hosts for Windows NLB, **223**
RBAC (role-based access control), **149–150**
RD Connection Broker server role
configuring, *453–458*, **453–458**
configuring RemoteApp, **459–465**, *459–465*
overview of, **423**
vs. role prior to Windows Server 2008 R2, 416
RD Gateway server role
configuring, **465–472**, *466–472*
overview of, **423–424**
vs. role prior to Windows Server 2008 R2, 416
RD Licensing Manager, 438, *438*
RD Licensing server role
configuring, **434–439**, *434–439*
overview of, **422–423**
vs. role prior to Windows Server 2008 R2, 416
RD Session Host (RDSH) server role
application diversity and, **421–422**
configuring, **439–449**, *441–449*
overview of, **417–418**
processor architecture, 420
usage profiles, **422**
virtualizing, 418
vs. role prior to Windows Server 2008 R2, 416
Windows architecture and, 418–419
RD Web Access server role
configuring, **450–452**, *450–452*
overview of, **424–425**
RemoteApp and, 417
vs. role prior to Windows Server 2008 R2, 416
RDMs (raw device mappings)
building failover clusters, *122–123*, 122–123, *126–127*, **126–127**
configuring VDR backup destination, **58–59**
mapping to VDR virtual machine, *60*, **60–61**
performing backup job, 62
virtual machine file system vs., 269–270, **322–323**, *323*
RDP (Remote Desktop Protocol), RD Gateway role using, 423
RDS (Remote Desktop Services) VM
application diversity, **421–422**
checking application compatibility, **427–428**, *428*
defining user requirements, **428–429**
deployment considerations, **431**
DRS and HA considerations, **432–433**
initial creation of RDS VMs, **431–432**
inventories for deployment, **425–427**
ongoing provisioning strategy, **432**
overview of, 412–413
processor architecture, **420–421**
RD Connection Broker role, **423**
RD Gateway role, **423–424**
RD Licensing role, **422–423**
RD Session Host role, **417–418**
RD Web Access role, **424–425**
single-server vs. farm models for deploying, **425**
sizing deployment, **429–431**
summary, 485–486
usage profiles, **422**
virtualization considerations, **415–417**
Windows architecture (32-bit and 64-bit options), **418–419**
RDS (Remote Desktop Services) VM, proof of concept
configuring Active Directory VM, **434**
configuring RD Connection Broker role, *453–458*, **453–458**
configuring RD Gateway server role, **465–472**, *466–472*
configuring RD Licensing role, *434–439*, **434–439**
configuring RD Session Host role, **439–449**, *441–449*
configuring RD Web Access role, *450–452*, **450–452**
configuring RemoteApp functionality, **459–465**, *459–465*
configuring Windows Network Load Balancing, **472–480**, *473–480*
monitoring farm deployment, **484–485**
overview of, 433
testing farm deployment, **480–484**, *481–483*
RDS Application Analyzer, 428, *428*
RDS Best Practices Analyzer, 447, *448*
RDS Session Host Configuration tool, *446*, 446
RDW (relational data warehouse), database workload categories, 327
read-only domain controllers (RODCs), **143**, **162**
Ready to Complete screen
backup job, **62**, *63*
creating new virtual machine, **98–99**, *99*
Deploy OVF Template Wizard, **228**, *228*
deploying VM from template, **115**
installing VDR appliance, **54**, *54*
recovery. *See* backup and restore; disaster recovery
Recovery Point Objective (RPO), Data Recovery, **23**, 52
Recovery Time Objective (RTO), Data Recovery, **23**, 52
redundancy
adding to IIS with WNLB, **198**
data, **248–249**
messaging and, 241
regedit.exe, time synchronization, **107**
registry
configuring time synchronization, **107–108**
modifying Windows Time Service, *147*, **147**
using Windows Time Service, **146**

relational data warehouse (RDW), database workload categories, **327**
Remediate Wizard, **75–78**, *76–78*
Remote Desktop Load Simulation Tools, 483–484
Remote Desktop Protocol (RDP), RD Gateway role using, 423
Remote Desktop Services. *See* RDS (Remote Desktop Services) VM
Remote Desktop Virtualization Host (RDVH), 415–417, *417*
Remote Desktop Web Access Configuration tool, *463*, 463
remote devices, application compatibility issues in RDS, 427
Remote Management service, IIS
  connecting to remote server, *209–211*, **209–211**
  enabling in Windows Server 2008 in full installation, **207–209**, *207–209*
  installing, **181**
remote office/branch office (ROBO) solution, **143**
RemoteApp Manager, *460*, 460
RemoteApps
  accessing from web browsers. *See* RD Web Access server role
  configuring, *459–465*, **459–465**
  providing user access to applications, 417
  testing RDS farm deployment, 480–483
RepAdmin (Replication Administration)
  characterizing Active Directory workload, 140
  monitoring Active Directory replication, **150–151**, 151, *173–174*, **173–174**
replication
  Active Directory VM network configuration and, **142**
  monitoring Active Directory VMs, **150–154**, *151–154*, *173–176*, **173–176**
ReplMon (Replication Monitor)
  characterizing Active Directory workload, 140
  monitoring Active Directory replication, **150–151**, *151*, **174**, *175*
  for Windows Server 2003, 154
reports
  failover cluster configuration, *132*, **132**
  monitoring status of backup job, **64**, *64*
request filtering, IIS 7 security, **198**
Requirements Calculator, *269*
Resource Authorization Policy (RAP), RD Gateway and, 424
resource management
  across multiple ESX servers with DRS, *8–9*, **8–10**
  improving scalability with DRS, **26**
  overview of, **6–7**
resource pools, DRS monitoring, *8*, **8–10**
resource requirements, sizing RDS deployment, 430
resource utilization baseline, existing SQL Server databases, **328–331**
Restart Policy setting, disabling for SQL Server clusters, **352**, *352*
restore. *See* backup and restore
`resxtop`. *See also* `esxtop`
  monitoring RDS virtual farms, 484
  performance counters, 485
RFC3484, **196**
ROBO (remote office/branch office) solution, **143**
role inventory, RDS deployment and, 425
role-based access control (RBAC), **149–150**
roles. *See also* FSMO roles
  CAS. *See* CAS (Client Access server) role
  installing Active Directory service, *160*, **160**
  installing IIS, *202*, **202**
  primary domain controller emulator, **144–146**, *145*
  SharePoint Server 2007 VM, **374–375**
rollback options, remediation, **77**, *77*
RPC client access, CAS server, 256
RPC encryption, CAS server, 241
RPO (Recovery Point Objective), Data Recovery, **23**, 52
RTO (Recovery Time Objective), Data Recovery, **23**, 52
rules, DRS resource allocation, **8–9**, *9*
Run Once, guest OS deployment, *113*, **113**
RVI (Rapid Virtualization Indexing), AMD, **5**, **44**, 420

S

SAM (standby Active Manager), Exchange Server, 249
SANs (Storage Area Networks), **23**, 277
SAS (serial attached SCSI)
  virtual device for Microsoft Cluster Service
    creating new VM, **95**, *96*
    as new generation of virtual hardware, **7**
    support for failover clustering, **23**
    virtual hardware version 7 using, 93
scalability services
  Distributed Resource Scheduler, **26**
  ESX and ESXi, **26**
  hot extend for virtual disks, **12**, **26**
  hot plug devices, **7**
  hot-add support, **26–27**, *27*
  overview of, **25**
  vNetwork, **18–19**
  vStorage, **14–15**
scale-out deployment strategy
  IIS performance, **183**
  overview of, **81**
  RDS, 431
  SharePoint database back-end server role and, 378
  SQL Server VMs, 333–334
scale-up deployment strategy
  IIS performance, **183**
  overview of, **81**
  SQL Server VMs, 334
scanning selected vSphere objects, *74–75*, **74–75**
SCC (single copy cluster), Exchange Server, 248–249
Schedule Update Download Wizard, Update Manager, *68*, **68**
schedules, domain controller backups, **170**, *170*
SCP (Secure Copy), copying ISO image, **100**
scripts, Exchange Server 2010, *288*
`Scripts` directory, installing Exchange Server prerequisites, *287*

SCSI controllers
building failover cluster, *124–125*, **124–126**
creating new VM, **95**, *96*
installing Windows 2008 R2, *103*
virtual hardware version 7 using, 93
Secure Sockets Layer (SSL), RD Gateway server role, 468
security
Active Directory time synchronization, **144**
domain controller virtual machine files, **149–150**
Exchange Server 2010 groups, *284*
IIS 7 features, **181**
planning virtual deployment, 45
read-only domain controllers for, **143**
SQL Server VMs and, **320**
strong passwords for, **157**, *158*
virtualized IIS, **198–201**, *200–201*
security ID (SID), guest OS deployment, **114**
security services
features, **24**
VMkernel, **25**
VMSafe, **24–25**
vShield Zones, **24**
Select and Configure a Raw LUN page, failover clusters, **122–123**, *123*
Select Compatibility Mode, failover clusters, **123**, *124*
Select Role Services screen, Add Roles Wizard for IIS, **204**, *204*
Select Web Server screen, Add Roles Wizard for IIS, **203**, *203*
serial attached SCSI. *See* SAS (serial attached SCSI) virtual device for Microsoft Cluster Service
Server Core
enabling remote management, **181**
as host for IIS installation, 180
Server Manager
configuring ESX for Windows NLB, **215**, *215*
installing for Active Directory domain controller backups, 169
installing IIS, *202*, **202**
preventing Active Directory deletions, 171–172, *171–172*
scanning Active Directory Domain Services, **165**
Server Virtualization Validation Program. *See* SVVP (Server Virtualization Validation Program)
service availability, **247–248**
Service Center Operations Manager
performance monitoring with, 272
planning virtual deployment, 46
Service Console, ESX and ESXi scalability, **26**
service-level agreements. *See* SLAs (service-level agreements)
`set-transportconfig`, 260
Setup program, Exchange Server 2010, *290*
Setup Support Rules, SQL cluster installation, *356*, 356
SharePoint, planning virtual deployment, **44**
SharePoint governance
areas of a governance plan, 372–374
overview of, 371–372
what it is, 372
SharePoint Server 2007 VM
application server role, **377**
database back-end server role, **378**
HP tool for sizing server roles, *386–395*, **386–395**
index server role, **377**
overview of. *See* SharePoint Server 2007 VM
roles and responsibilities, **374–375**
SharePoint governance, **371–374**, *372*
sizing server roles, 378–386, **379–386**
summary, 413–414
web front-end server role, **376–377**
workloads and, **375–376**
SharePoint Server 2007 VM, proof of concept
building/configuring database back-end VM, **396–402**, *397–402*
configuring Central Site Administration VM, **402–404**, *403–404*
configuring index server VM, **404–405**, *405*
configuring WFE server VM, **405–406**, *406*
configuring Windows Network Load Balancing, **409–412**, *410–412*
installing SharePoint software, **406–409**, *407–409*
overview of, 395–396
SharePoint Server 2010, features in, 412–413
Shavlik patch management, Update Manager integration with, **34**, **65**
SID (security ID), guest OS deployment, **114**
single copy cluster (SCC), Exchange Server, 248–249
Site Recovery Manager (SRM), **36–38**, 324
Sizing Estimator, Update Manager, **66**
SLAs (service-level agreements)
backing up VMs using VDR, 52
characterizing Active Directory workload, 140
Data Recovery supporting, **23**
planning virtual deployment, 45
SMASH (System Management Architecture for Server Hardware) profiles, **32**
SMTP, EDGE server and, 238, 250
snapshots
converting physical machine to virtual machine, **116–118**
Data Recovery using, **23**
potential to corrupt AD domain controllers, 62, *148*, **148–149**
Remediate Wizard and, **77**
software
failover cluster requirements, **119**
firewalls, *200*, **200**
load balancing IIS server, **196**
migrating physical machine to virtual and, **117**
RDS deployment and, 426
user requirements in RDS, 429
vStorage improvements for NFS and iSCSI, **14–15**

spam controls
  EDGE server, 238
  Source Discovery, 277
SpecWeb 2005 workload, **183**
SQL Server
  2008 Enterprise edition, 398–399, *399*
  scale-out deployment strategy for, **81**
  VM placement in vSphere cluster and, **82**
SQL Server VM
  32-bit vs. 64-bit, **317–318**
  availability requirements, **323–325**, **331**
  backup considerations, **325–326**
  backup strategies and requirements, **331–332**
  capacity and system growth requirements, **332**
  configuring, **340–346**, *341–345*
  data integrity and, **314**
  database consolidation, *315*, **315–316**
  deployment strategy, **333–335**
  evaluating existing SQL servers for, 326–327
  hardware, **321**
  licensing, **317**
  migrating to virtual environment, **335–337**
  overview of, 313
  patching and update strategies, **332**
  performance, **316–317**, *317*
  ranking candidates for, **333**
  resource utilization baseline for existing databases, **328–331**
  security, **320**
  storage, 321–322
  summary review, 370
  support considerations, **317–320**, *319–320*
  virtualization considerations, **313–314**
  VMFS vs. RDM, **322–323**, *323*
  what not to virtualize, **326**
  workload evaluation, **327–328**
SQL Server VM, proof of concept
  Cluster Installation Rules, **357**, *357*
  configuring, **340–346**, *341–345*
  configuring failover cluster VMs, 350–351
  configuring virtual disks, **348–350**, *349–350*
  configuring vSphere, **338–340**
  creating and validating failover clusters, **353–354**, *353–354*
  disabling VM restart priority, *352*, **352**
  DRS anti-affinity rules, *351–352*, **351–352**
  first node of failover cluster installed, *357*, **357**
  installing SQL Server in failover cluster, **355–356**, *355–356*
  large memory pages used for, **347**
  locking buffer pool pages in memory, *346*, **346–347**
  Microsoft Distributed Transaction Coordinator and, *354*, **354**
  overview of, 337–338
  performance monitoring, **362–370**, *363–367*, *369*
  second node of failover cluster installed, **358–359**, *358–359*
  selecting disk resources for failover cluster, **356**, *356*
  testing deployment, **360–362**, *361–362*
  testing failover cluster services, **359–360**, *360*
SQLIOSim tool, 360–361, *361*, 366
SRM (Site Recovery Manager), **36–38**, 324
SSL (Secure Sockets Layer), RD Gateway server role, 468
Standard Switch, vNetwork, **15–17**
Standard version of Windows Server 2008, for IIS VM, 180
standby Active Manager (SAM), Exchange Server, 249
storage. *See also* vStorage services
  in Client Access server role, 257–258
  in Edge Transport server role, 252–255
  for Mailbox server role, 267–270
  performance and scalability, **14–15**
  performance monitoring for UM servers, 273
  RDS Session Host servers and, 440
  SQL Server VM considerations, 321–322
  in Unified Messaging server role, 271–272
Storage Area Networks (SANs), **23**, 277
Storage VMotion
  defined, **20**
  overview of, **21**, *22*
  VMFS as key enabler for, **11**
stratum, **144**
strong password requirement, **157**, *158*
Summary tab, viewing VM configuration in, **100**
support, SQL Server VM considerations, **317–320**, *319–320*
Support Clustering Features Such As Fault Tolerance option
  building VM golden image, **97**, *98*
  creating virtual disks, **119**
Support Policy Wizard, SVVP, *79–80*, **79–80**, *319–320*, **319–320**
SVGA driver, Tools package installer not installing, **105**
SVVP (Server Virtualization Validation Program)
  deployment strategies, **81–82**
  history of, 236
  new era of virtualization and, 42
  overview of, **78–81**, 78–81
  working with, *319–320*, **319–320**
switches, vNetwork, **16–18**
synchronization, time. *See* time synchronization
Sysprep, **110**, **116**
System Center Capacity Planner
  application configuration page, *383*
  customizing workload transaction mix, *381*
  hardware selection, *383*
  Model Wizard screen, *381*
  Networks configuration page, *382*
  overview of, 379
  selecting SharePoint Server in drop-down, *380*
  simulation results page, *385*
  sizing SharePoint server roles, 379–386, *380–386*
  summary page, *384*
  visually representing SharePoint topology in, *386*
System Management Architecture for Server Hardware (SMASH), **32**

T

teams and roles area, SharePoint governance plan, 373
templates, VM
  creating multiple VMs, **49**
  creating RDS VM, 432
  creating SQL Server VM, 335–336
  creating Windows Server 2008 VM, **88**
  overview of, *108–109*, **108–110**
  procedure for, **114–116**, *115–116*
  readying guest OS, **110–114**, *111–114*
Terminal Services. *See also* RDS (Remote Desktop Services)
  VM, 415
test phase, virtual deployment, 46
testing
  deploying SQL Server VM, **360–362**, *361–362*
  Exchange Server solution, **298–302**, *299–302*
  failover cluster configuration, **131–132**, *132*
  failover cluster services, **359–360**, *360*
  RDS farm deployment, **480–484**, *481–483*
text-to-search, UM server role, 245
thick-provisioned disk, **163**
thin provisioning. *See* virtual disk thin provisioning, vStorage
third-party applications
  backup VMs, **136**
  impacting mailbox performance, 268
  load-balancing appliances, **196**
  vDS switch support for, **18**
Threat Management Gateway (TMG), **200**
tier-1 applications
  new era of virtualization and, 41–42
  performing backup of, **62–65**, *63–65*
time synchronization
  controlling time drift, **144**, *145*
  deploying Windows 2008 VM, **107–108**
  example of improper, **165**, *165*
  implementing Active Directory VM, *147*
  Kerberos v5 for, **144**
  modifying Windows Time Service, **147**, *147*
  monitoring Active Directory replication, **150–154**, *151–154*
  proof of concept, *166–168*, **166–168**
  using Group Policy and WMI filters, **146**
  using Windows Time Service, **146**
  in virtual machines, **107–108**
timekeeping, in virtual machines, **107–108**
TLB (translation lookaside buffer), memory management, **5**, **6**
TMG (Threat Management Gateway), **200**
tombstone lifetime setting, system-state or domain controller backups, **170**
Tools, VMware
  disabling time synchronization, **107**
  installing in Server Core edition of Windows 2008/R2, **105**
  installing on server, *158*, **158–159**
  installing on Windows 2008 R2, **104–105**, *105*
  time synchronization, **147**
TPC (Transaction Processing Council), 362
TPM (Trusted Platform Module), VMkernel, **25**
tracking, Chargeback automation of, **36–37**
traffic-shaping policies, vDS, **18**
transaction logs, SQL Server VM, 342, *342*
Transaction Processing Council (TPC), 362
translation lookaside buffer (TLB), memory management, **5**, **6**
transparent page sharing, memory management, **6**
transport agents, EDGE server, 251
transport dumpster, Hub Transport server, **258–261**
Trusted Platform Module (TPM), VMkernel, **25**
tuning performance, IIS VMs, **189–195**, *191–195*
type 1 virtualization, **3**, **43**
type 2 (or hosted hypervisor) virtualization, **3**, *42*, **42–43**

U

UM (Unified Messaging) server role
  dual-homed UM virtual machine, *270*
  functions of, 270–271
  monitoring performance of, 272–274
  overview of, **244–245**
  processor configurations, 271
  storage in, 271–272
unicast mode, Windows Network Load Balancing, **197**
Update Manager
  attaching baselines/baseline groups to vSphere objects, *74*, **74**
  configuring patch download source, *67*, **67**
  creating baselines/baseline groups, *68–73*, **68–73**
  defined, **27**
  deployment, *65*, **65–66**
  downloading patches, extensions and related metadata, *68*, **68**
  overview of, **34**, 65
  remediating selected vSphere objects, **75–78**, *76–78*
  scanning selected vSphere objects, **74–75**, *74–75*
  Sizing Estimator, **66**
  vSphere, 432
updates
  SQL Server VM, **332**
  using vSphere Host Update Utility, **35**
upgrade baselines, *72*, **72**, **74**
upgrades
  using Update Manager, **34**
  VM hardware, **7–8**, *7–8*
uptime and downtime area, SharePoint governance plan, 373–374
usage profiles, RD Session Host role and, **422**
user groups, RD Gateway server role, 469, *469*
users
  placing domain controllers close to, **143**
  RDS deployment and, 426
  requirements in RDS, **428–429**, 429
USN numbers, Active Directory replication, **150**

V

validation, SQL Server clusters, **353–354**, *353–354*
validation wizard, failover cluster configuration, **131–132**, *132*
vApps, **35**
VCB (VMware Consolidated Backup), 325
vCenter
  AppSpeed, **35–36**
  AppSpeed tool, *367*
  Capacity IQ, **35–36**
  Chargeback, **36–37**
  Converter. *See* Converter
  customizing for performance monitoring, 366–367, *367*
  Guided Consolidation, **27**, **33–34**
  Lab Manager, **36**, **37**
  Lifecycle Manager, **36**, **37**
  Maps view, **142**
  monitoring RDS virtual farms, 484
  monitoring SQL Server VMs, *363–365*, 363–365
  Orchestrator, **27**, **33**
  performance counters, 485
  Site Recovery Manager, **36–38**
  Update Manager. *See* Update Manager
  Update Manager Sizing Estimator, **66**
vCenter Server
  centralized licensing, **33**
  defined, **27**
  editions, **39**
  hardware monitoring with CIM, *32*, **32–33**
  host profiles, **29–30**, *30–31*
  linked mode, **31–32**
  management services. *See* management services
  navigation and inventory search, **28–29**, *29*
  overview of, **28**
  working with VMFS Volume Grow, **12**
vCLI (vSphere CLI), **35**
vCompute services
  Distributed Power Management, **2**, **10**, *10*, **31**
  Distributed Resource Scheduler. *See* DRS (Distributed Resource Scheduler)
  ESX and ESXi. *See* ESX/ESXi servers
  overview of, *2*, **2**
VDI (virtual desktop infrastructure), 415, **416–417**, *417*
VDR (Data Recovery)
  configuring appliance, **55–62**, *55–62*
  defined, **20**
  installing appliance, **52–54**, *53–54*
  overview of, **23**, **52**
  performing backup of tier-1 application, **62–65**, *63–65*
  VM backups with, **135–136**, *136*
  vSphere backup services, 325
VDR Getting Starting Wizard, *58*, **58**
vDS (Distributed Switch), vNetwork, **17–18**
VEMs (virtual Ethernet modules), and vDS third-party switch support, **18**
version, specifying VM, *93*, **93**
VI3 (Virtual Infrastructure 3), ESX CPU scheduler in, **5**
video functions, Tools for, **158**
viewing templates, *109*, **109**
Virtual Appliance Marketplace, **196**
virtual desktop infrastructure (VDI), 415, **416–417**, *417*
virtual disk thin provisioning, vStorage
  defined, **11**
  overview of, **12**, *13*
  selecting for new VM, **97**, *98*
  using Storage vMotion with, **21–22**, *22*
virtual Ethernet modules (VEMs), and vDS third-party switch support, **18**
virtual hardware generation, ESX and ESXi new, **7–8**, *7–8*
Virtual Infrastructure 3 (VI3), ESX CPU scheduler in, **5**
virtual machine disk files. *See* VMDK (virtual machine disk) files
Virtual Machine File System. *See* VMFS (Virtual Machine File System)
virtual machine files, securing domain controller, **149–150**
Virtual Machine Interface (VMI), **4**
Virtual Machine Properties page, configuring VDR, **60**, *60*
Virtual Machine Versions page, Create New Virtual Machine Wizard, **93**
virtual memory address translation, hardware-assisted, **4–5**
virtual network adapters. *See also* VMXNET3 network adapter
  Active Directory VM network configuration and, **142**
  defined, **15**
  overview of, *16*, **16**
virtual storage adapters. *See* PVSCSI (paravirtualized SCSI) storage adapter
virtual supervisor modules (VSMs), **18**
virtualization architectures, **3**
virus controls
  EDGE server, 238
  Source Discovery, 277
  VMSafe, **25**
vLockstep technology, Fault Tolerance using, **118**
VM Network
  cluster node build specification, 120
  creating VMs, 95
  failover cluster configuration, **133**
VMDirectPath I/O for Networking, **19**
VMDirectPath I/O for storage, **15**
VMDK (virtual machine disk) files
  building failover cluster, **123**, *123*, *126*, **126**
  installing AppSpeed using, 189
  location, **97–98**
  selecting for new VM, **97**, *98*
  SharePoint disk requirements, 377
VMFS (Virtual Machine File System)
  copying ISO image to, **100**
  enabling VMotion, **20**
  hot extend for, **12**
  overview of, **11**
  preventing I/O performance degradation, 343
  SharePoint disk requirements, 377
  SQL Server VM storage considerations, 321
  VMFS Volume Grow, **12**
  vs. RDM, 269–270, **322–323**
VMI (Virtual Machine Interface), **4**
VMkernel
  ESX and ESXi scalability with 64-bit versions of, **26**
  security service, **25**
  Storage VMotion migration using, **21**
  vmknic network adapter for, **16**
vmknic network adapter, **16**, **119**

VMotion
administering with vCenter Server, **28**
availability service, **20–21**, *21*
defined, **19–20**
licensing, **47**
relocating SQL Server VMs to other ESX servers, 332
VMFS as key enabler for, **11**
VMSafe, **24–25**
VMStatsProvider, **186**
VMTX files, **108**
VMware Consolidated Backup (VCB), 325
VMware File System. *See* VMFS (Virtual Machine File System)
VMX files, virtual machines, **108**
VMXNET3 network adapter
building failover cluster, **121**, **126**
configuring ESX for Windows NLB, **212**, *212*
defined, **15**
features of, *16*, **16**
as new generation of virtual hardware, **7**
performance enhancement of, **18**
PVSCI storage adapter vs., **13**
virtual hardware version 7 using, 93
VMware Tools for, **158**
vNetwork
Distributed Switch, **17–18**
features, **15**
overview of, *2*
performance and scalability, **18–19**
Standard Switch, **16–17**
virtual network adapter. *See* VMXNET3 network adapter
VMDirectPath I/O for Networking, **19**
voice communication, UM servers and, 270
Voice over IP (VoIP), UM servers and, 270
VoIP (Voice over IP), UM servers and, 270
volume manager, VMFS acting as logical, **11**
Volume Shadow Copy Service. *See* VSS (Volume Shadow Copy Service)
`vscsiStats`, SQL Server VMs, 367–369
vShield zones, **24**, **200–201**
VSMs (virtual supervisor modules), **18**
vSphere
application services. *See* application services
command-line interfaces, **28**, **35**
compatibility guides, **39–40**
configuring for SQL Server VMs, *338–340*, **338–340**
DRS (Distributed Resource Scheduler). *See* DRS (Distributed Resource Scheduler)
editions, **38–39**
HA (high-availability). *See* HA (high-availability)
Host Update Utility, **35**
infrastructure services. *See* infrastructure services
Management Assistant (vMA), **34**
management services. *See* management services
number of SQL Server VMs per host, 335
performance counters, 365, 485
PowerCLI, **35**
provisioning RDS VMs, 432
storage improvements for addressing performance issues, 317
summary review, 39–40
VMFS vs. RDM, 323
VSS (Volume Shadow Copy Service)
converting physical machine to virtual machine, **116–117**
Data Recovery supporting, **23**
Data Recovery using, **135–136**
vSphere backup services, 325
Windows Server Backup, **135–136**
vStorage services
APIs, **14**
key features of, **11**
overview of, *2*
paravirtualized SCSI. *See* PVSCSI (paravirtualized SCSI) storage adapter
planning virtual deployment, 45
storage performance and scalability, **14–15**
thin provisioning, **12**, *13*
Virtual Machine File System. *See* VMFS (Virtual Machine File System)
VMDirectPath I/O, **15**
VMFS Volume Grow, **12**
vswif network adapter, **16**
vSwitch, deploying clustered VMs, **119**
VT-x (Virtualization Technology), Intel, **44**

W

web browsers. *See also* RD Web Access server role, 450
Web Server (IIS) screen, Add Roles Wizard for IIS, *204*, **204**
Web version of Windows Server 2008, choosing for IIS VM, 180
welcome screen, configuring VDR, 56
WFC (Windows Failover Clustering). *See also*
failover clustering
building, **120–127**, *122–127*
configuring, **127–134**, *127–134*
configuring cluster for database back-end VMs, 397
creating failover clusters, 353
deploying, **118–119**
failover cluster management with, *355*
hardware/software requirements, **119**
high-availability solutions in Exchange Server, 249
limitations of in vSphere, **120**
overview of, **118**
setting up for SQL Server VMs, 347
support for, **23**
WFC (Windows Failover Clustering), deploying
building Microsoft failover cluster, **120–127**, *122–127*
clustering VMs, **118–119**
configuring, *127–134*, **127–134**
hardware/software requirements, **119–120**
overview of, **118**
WFE (web front-end) server role, SharePoint Server
configuring, **405–406**, *406*
responsibilities of, 374–375
workload related to, **376–377**

Windows
architecture, 418, **418–419**
authentication in IIS 7, **199**
Perfmon. *See* Perfmon
specifying 7 as VM version, **93**
Windows Failover Clustering. *See* WFC (Windows Failover Clustering)
Windows Management Instrumentation (WMI) filter, time synchronization, **146**, **166–168**
Windows Network Load Balancing. *See* WNLB (Windows Network Load Balancing)
Windows Server 2003, **23**
Windows Server 2008
choosing flavor for IIS VM, **180**
high-availability clustering with, **23**
improvements to kernel memory area, 419
RDS roles in, 416
Windows Server 2008 R2
creating new VM for, **90–100**
installing, *101–106*, **101–106**
installing for Active Directory VMs, *156–158*, **156–158**
specifying as VM version, **93**
Windows Server 2008 VMs
deployment. *See* deployment, Windows 2008 VMs
guest OS compatibility and, **86–87**
guest OS requirements and, **87–88**
overview of, 85
protecting your data, **134–136**, *135–136*
summary review, 137
understanding guest OS compatibility and requirements, **86–87**
Windows Server Backup
domain controller backups, *169*, **169–171**
overview of, **134–135**, *135*
recovery of Active Directory VMs, **149**
Windows System Resource Manager (WSRM), *370*
Windows Time Service
modifying, **147**, *147*
synchronization parameters, **166–167**
time synchronization using, **146**
WINS (Windows Internet Name Service), 140
WMI (Windows Management Instrumentation) filter, time synchronization, **146**, **166–168**, *168*
WNLB (Windows Network Load Balancing)
configuring, **409–412**, *410–412*
configuring ESX for, *211–222*, **211–222**
configuring for RDS farms, **472–480**, *473–480*
virtualizing IIS, **196–198**
workflows, **33**
workgroup, guest OS deployment, **113**
workloads
evaluating SQL Servers for virtualization, **327–328**
implementing Active Directory VMs, **140–141**
planning virtual deployment, 46
VM placement in vSphere cluster, **82**
workloads, in Exchange Server 2010
of Client Access and Hub Transport combined server, **261–262**
of Client Access server, **255–258**
of Edge Transport server, **250–255**
of Hub Transport server, **258–261**
of Mailbox server, **262–270**
overview of, **250**
of Unified Messaging server, **270–274**
workloads, in SharePoint Server 2007
application server role, **377**
capacity planning, *381*
database back-end server role, **378**
index server role, **377**
overview of, **375–376**
WFE (web front-end) server role, **376–377**
Workstation, as hosted virtualization product, 43
WSRM (Windows System Resource Manager), *370*

X

x86 servers
binary translation virtualization for, **3**, **44**
infrastructure services virtualizing, 1–2, *2*
virtualization products for, **3**

Z

ZIP file, installing AppSpeed, **189**